

Some introductory notes on game theory

The mathematical analysis in the preceding chapters, for the most part, involves nothing more than algebra. The analysis does, however, appeal to a game-theoretic vocabulary and set of concepts that may be unfamiliar. This Appendix introduces those concepts and vocabulary in order to give readers with little or no background in game theory a better sense of the tools used to analyze deterrence theory and some of the strengths and weaknesses of those tools.¹

The extensive form

The brinkmanship and limited-retaliation models are examples of games in extensive form. A *game in extensive form* is composed of two parts.² The first is the *game form* or *game tree*. The second is the players' *payoffs*. The game form or tree is an abstract summary of the situation facing the players. The tree tells the order of play, the set of alternatives from which each player must choose when it plays, and what each player knows when it must choose. The tree defines who moves after whom, what each player can do, and what each player knows about what the other players have done when it must decide what to do.

Two very simple trees are illustrated in Figure A1. In both Figure A1(a) and A1(b) the order of play is the same. Player *I* moves first, and then player *II* moves. When *I* moves, the trees show that it can choose between two alternatives: It can choose up, *U*, or down, *D*. Similarly, *II* has only two alternatives: top, *T*, and bottom, *B*. The trees also define what *II* knows when it must decide between *T* and *B*. In Figure A1(a), *II* is assumed to know what *I* did, perhaps because *II* could watch *I*. In Figure A1(b), however, *II* does not know what *I* did. This is the meaning of the dashed line connecting *II*'s two decision nodes in Figure A1(b). Of course, *II* may have beliefs about whether it is at its upper or lower decision node, and more will be said later about beliefs and their formation. At this point, it is important

¹ For an excellent though somewhat more technical introduction to game theory than the one presented here, see Tirole (1988, pp. 423–59).

² For a formal definition, see Luce and Raiffa (1957), Owen (1982), Selten (1975), or Kreps and Wilson (1982b).

to note that only the tree in Figure A1(b) is intended as a model of a situation in which *II* must decide what to do without knowing what *I* has done. There is simply not enough information.


If a player is unable to distinguish between some of its decision nodes, then these indistinguishable nodes constitute an *information set*. In Figure A1(b), *II* has one information set because its two decision nodes are indistinguishable. But in Figure A1(a), *II* can distinguish between its two decision nodes because it knows what *I* did when it has to decide what to do. *II*, therefore, has two information sets in this tree, each composed of a single node. In both Figure A1(a) and A1(b), *I* has a single information set consisting of a single decision node.

If, as in Figure A1(a), every information set consists of a single decision node or *singleton*, so that a player at any information set knows exactly what alternatives the other players have previously played, then the game has *perfect information*. Chess is a game of perfect information. Whenever a player must decide what to do in chess, it is completely certain of what all of the preceding moves have been. That is not the case in the tree in Figure A1(b), where *II* does not know what *I* has done.

The *terminal nodes* of a tree are the points at which a path through the tree ends. In Figure A1(a) and A1(b), for example, there are four terminal nodes, each of which follows one of the four branches that *II*'s decision can take. The terminal nodes correspond to the possible *outcomes* of the game.

The game tree abstractly defines the situation in which the players must act. Each path through the tree leads to a terminal node that is associated

Figure A1. Some simple game trees.



with some possible outcome. But to have any hope of analyzing what will be done in this situation, more than the structure of the situation must be described; the players' preferences over the possible outcomes must also be defined. This is the second part needed to complete the specification of the game. That is, each player associates with each possible outcome a *payoff* or *utility* that reflects its preferences over the set of possible outcomes.³

Three examples will make this description more concrete. The first is the game of chicken. There are two players: *I* and *II*. Each player has two alternatives: It can stand firm, *F*, or submit, *S*. The decision whether to stand firm or submit is made in ignorance of what the other player is doing. The game tree in Figure A2(a) illustrates this situation. The tree begins with *I* having to decide between *F* and *S*. *II* must then decide between *F* and *S*. The tree also models the assumption that neither player knows what the other player is doing when it must decide what to do. *I* clearly cannot determine what *II* has done when *I* is making its decision, because the information set at which this decision is made precedes the information set at which *II* makes its decision. Similarly, *II* does not know what *I* is doing because both of *II*'s decision nodes are in the same information set, and, by definition, a player cannot distinguish among the nodes in any one of its information sets. *II* cannot tell if it is at its upper node, in which case *I* is standing firm, or if it is at its lower node, in which case *I* is submitting. In the tree, both players make informationally isolated decisions; each player must decide what to do without knowing what the other player is doing. One natural interpretation of this informational isolation is that the tree is a model of a situation in which decisions are made simultaneously. That is, in the actual situation for which Figure A2(a) is a model, *I* and *II* make their decisions to stand firm or submit simultaneously. Simultaneity, in turn, implies that no player can know what the other is doing when it must decide what to do. In this way, simultaneity makes for informationally isolated decisions, and that is what is modeled in the tree in Figure A2(a).

To complete the specification of the game of chicken, the players' payoffs or preferences over the possible outcomes of the game must be specified. If one player stands firm and the other submits, then the player who stands firm wins, and the other loses. If both stand firm, there is a disaster that is worse than losing. If both submit, a compromise results that is better than losing, but not as good as prevailing. Picking numbers to represent these payoffs, suppose that if one player stands firm and the other submits, then the player who stands firm receives 1, whereas the player who submits loses

³ Usually, utilities are assumed to be von Neumann–Morgenstern utilities. That is, the utility of an uncertain event is the expected utility of the possible events. For example, the utility of a lottery that will give utility u_1 with probability ρ and utility u_2 with probability $1 - \rho$ is $\rho u_1 + (1 - \rho)u_2$.

1. If both players stand firm, both lose 5. If both submit, then each obtains the compromise payoff of zero. Thus, the payoffs at the end of the branch along which *I* plays *F* and *II* plays *S* are $(1, -1)$, where the first element in the pair of payoffs is *I*'s payoff, and the second is *II*'s. The complete specification of the game is given in Figure A2(b).

The second example is the game of matching pennies. In this game, two players act simultaneously, and each reveals one side of a penny. If both players show heads or tails, player *I* wins and collects a penny from *II*. If one player shows heads and the other tails, then *II* wins and takes a penny from *I*. The extensive form of this game is depicted in Figure A3. *I* begins by making an informationally isolated decision between heads, *H*, and tails, *T*, after which *II* makes an informationally isolated decision between *H* and *T*. If *I* and *II* play the same face, the payoffs are $(1, -1)$ and if they make different choices, the payoffs are $(-1, 1)$.

Finally, consider a more complicated game that is a much-simplified version of poker. In this game, one card is dealt to player *I*, and another card is dealt to *II*. Each player can see only the card dealt to it. Then, knowing its card, but not its opponent's, *I* must decide whether to bid a dollar, *B*, or fold, *F*. If *I* folds, it loses its ante of one dollar to *II*. If *I* bids, *II* must either bid a dollar or fold. If *II* folds, *I* collects *II*'s ante of a dollar. If *II* bids, then both players expose their cards. If both cards are of the same color, the players divide the pot, which leaves a net gain of zero. If the colors differ, then black beats red, and the player holding the black card collects the pot of four dollars for a net gain of two dollars.

Figure A2. Chicken in extensive form.




Figure A3. Matching pennies.





Figure A4 shows the extensive form of this simple poker game. A player called "Nature" or *N* makes the first move. Assuming there to be a player called Nature is simply a modeling device used to introduce random or probabilistic elements into the game. For example, four combinations of colored cards could be dealt in the game, (B, B) , (B, R) , (R, B) , (R, R) , where the first element of the pair corresponds to the color of *I*'s card, and the second element is the color of *II*'s card. To represent this in the game, Nature begins the game by playing one of the four alternatives, where each alternative corresponds to one possible deal. With a very large deck, the probability of dealing one of these combinations will be $\frac{1}{4}$, and so Nature will play each of these alternatives with probability $\frac{1}{4}$.

After *N* plays, *I* must decide whether to bid or fold. When making this decision, *I* knows the color of its card, but not the color of its opponent's card. Accordingly, *I* cannot distinguish between a deal of (B, B) and (B, R) or between a deal of (R, B) and (R, R) . This means that *I* has two information sets, with the nodes representing the deals (B, B) and (B, R) in one information set, and the nodes representing the deals (R, B) and (R, R) in the other. At these information sets, *I* has two alternatives: bidding, *B*, or folding, *F*. If it folds, the game ends, and the payoffs are $(-1, 1)$. If *I* bids, *II* must then decide whether to bid or fold. *II*, like *I*, knows only the color of its card and consequently cannot distinguish between the deals of (B, B) and (R, B) or between the deals of (B, R) and (R, R) . *II*, therefore, also has two information sets, as shown in Figure A4. If *II* folds, the payoffs are $(1, -1)$. If *II* bids, the players expose their cards and obtain the payoffs described earlier and illustrated in Figure A4.

Strategies and the normal form

Now that a game has been described, one can begin to discuss ways of analyzing it. The first step is to define what is meant by a player's *behavioral strategy*. A player's behavioral strategy is simply a complete plan for how this player will play the game. This strategy tells what this player will do in each contingency that might arise in the game. More formally, a player's behavioral strategy is a rule that specifies which alternative this player will select at each of its information sets. If, as in the games of chicken or matching pennies, a player has a single information set, then a player's behavioral strategy merely tells what this player will do at this one information set. In matching pennies, a strategy for *I* is to play *H*. A second strategy for *I* would be to play *T*. In the simple poker game in Figure A4, each player has two information sets. Accordingly, a player's behavioral strategy must specify what the player will do at both of its information sets. A behavioral strategy for *I* is "fold if *I*'s card is red, and bid if the card is black." The instruction "fold if *I*'s card is red" cannot be a behavioral strategy for *I*, for it is not a complete plan for playing the game; it does not specify what *I* is to do if it is dealt a black.

Figure A4. The extensive form of the simple poker game.



It will be useful to distinguish between pure behavioral strategies and mixed behavioral strategies. In a *pure strategy*, the rule defining a player's strategy specifies that the player is certain to choose a single alternative at each of its information sets. In a *mixed strategy*, a player is allowed to randomize over the alternatives from which it must choose. That is, the rule defining a player's mixed behavioral strategy specifies a probability distribution over the set of alternatives at each of this player's information sets. This distribution gives the probability that any of the alternatives available at a given information set will be played. In matching pennies, for example, *I* has two pure strategies. It can play *H* for sure, or it can play *T* for sure. A mixed behavioral strategy for *I* in this game would be to show heads with probability $\frac{1}{2}$ and tails with probability $\frac{1}{2}$. A second mixed behavioral strategy would be to play *H* with probability $\frac{2}{3}$ and *T* with probability $\frac{1}{3}$. In the simple poker game, a pure behavioral strategy for *I* would be to "bid regardless of the color of *I*'s card." A mixed behavioral strategy would be to "bid with probability $\frac{3}{4}$ and fold with probability $\frac{1}{4}$ if *I*'s card is black, and bid with probability $\frac{3}{8}$ and fold with probability $\frac{5}{8}$ if *I*'s card is red."

The use of pure behavioral strategies makes it possible to define the *normal form* of a game, which is more compact and sometimes more useful in analyzing the game than the extensive form. Suppose that there are *M* players in some extensive-form game. Let s_i be some pure behavioral strategy for player *i*. That is, s_i is a rule that tells which alternative *i* is certain to play at each of its information sets. Now consider the *m*-tuple (s_1, s_2, \dots, s_M) , where each s_i in this *m*-tuple is a complete plan for how player *i* will play the game. This means that (s_1, s_2, \dots, s_M) describes what will be done at every information set in the game. Accordingly, one can imagine giving the plan (s_1, s_2, \dots, s_M) to a referee and then having the referee play out the game according to the players' strategies. If, for example, the first information set in the tree belonged to player *i*, then the referee would consult s_i in (s_1, s_2, \dots, s_M) to see which alternative *i* would choose at that information set. The referee would then follow the branch in the tree corresponding to that alternative and go down the tree to the next information set. If that set belonged to *k*, the referee would consult s_k to see what s_k would have *k* do at that information set. In that way, the referee could follow the plan defined by (s_1, s_2, \dots, s_M) and eventually reach a terminal node that would mark the end of the game. Put another way, the plan (s_1, s_2, \dots, s_M) defines a path through the tree, or, if Nature is making random moves in the tree, the plan defines the probability of reaching each possible terminal node or outcome of the game. Now recall that each player attaches some utility to every possible outcome of the game. If, therefore, a plan like (s_1, s_2, \dots, s_M) defines the probabilities of reaching the possible outcomes, then each player can attach an expected utility to the plan. That

is, each player knows what its expected utility will be if the game is played according to the plan (s_1, s_2, \dots, s_M) . Let $U_i(s_1, s_2, \dots, s_M)$ be the utility player i receives if the game is played according to (s_1, s_2, \dots, s_M) . Now, the game can be described by the set of all possible plans (i.e., the set of all possible m -tuples of pure behavioral strategies, and the utility functions that specify the utility the players receive if the game is played according to a specific plan). This description is the normal form of the game.

To make this description of the normal form more concrete, the extensive-form representations of the games of chicken, matching pennies, and poker will be translated into their normal forms. In chicken and matching pennies, each player has two pure behavioral strategies. This means that there are four different plans for playing the game. One way to keep track of these plans is with a matrix, where each row corresponds to one of player I 's strategies, and each column corresponds to one of II 's strategies. Each cell in the matrix then corresponds to a different combination of I 's and II 's strategies or, in other words, to a different complete plan for playing the game. The utility each player receives if the game is played according to a particular plan is placed in the cell associated with that plan. The normal form for chicken is shown in Figure A5(a), and that for matching pennies in Figure A5(b).

To translate the simple poker game into its normal form, note that I has four pure strategies. One strategy is to bid if a black card is dealt and to fold if a red card is dealt. Let $\{(B, b), (R, f)\}$ denote this strategy, where the first element in a parenthetical pair stands for the color of the card that may be dealt, and the second element tells what to do if this color is actually drawn; so (B, b) means bid, b , if a black card, B , is dealt. Then the other three strategies are $\{(B, b), (R, b)\}$, $\{(B, f), (R, b)\}$, and $\{(B, f), (R, f)\}$. Player II also has the same four strategies. It can bid or fold depending on whether it has a red or black card. Because each player has four strategies, there are sixteen different combinations of strategies (i.e., sixteen different complete plans for

Figure A5. The normal forms of chicken and matching pennies.

		II	
		F	S
I	F	-5, -5	1, -1
	S	-1, 1	0, 0

(a)

		II	
		H	T
I	H	1, -1	-1, 1
	T	-1, 1	1, -1

(b)

playing the game). As before, one can keep track of these different combinations in a matrix, where each row corresponds to one of I 's strategies, and each column corresponds to one of II 's strategies. This is done in Figure A6.

As an example of how the payoffs are calculated, consider the cell associated with I 's strategy of bidding if it has a black card and folding if it has a red card, which is denoted $\{(B, b), (R, f)\}$, and with II 's strategy of folding with a black card and bidding with a red card, which is given by $\{(B, f), (R, b)\}$. This corresponds to the cell at the intersection of the second row and the third column, where the payoffs are $(\frac{3}{4}, -\frac{1}{4})$. To derive these payoffs, suppose that Nature deals a red card to I and a red to II ; then play follows the branch (R, R) in the extensive form in Figure A4. Given that I is holding a red card, its strategy is to fold. The game ends with payoffs $(-1, 1)$. Now suppose that Nature deals a black card to I and a red card to II . Play then proceeds down the (B, R) branch. I 's strategy is to bid. Because II is holding a red card, it also bids. Because black beats red, the payoffs are $(2, -2)$. If Nature had dealt two black cards, I would have bid, but II would have folded, leaving the players with $(1, -1)$. Finally, a deal of red to I and black to II has I folding immediately, to give the payoffs $(-1, 1)$. Because Nature will deal each of these combinations with probability $\frac{1}{4}$, the expected

Figure A6. The simple poker game in normal form.

		II			
		$\{(B, b), (R, f)\}$		$\{(B, f), (R, f)\}$	
		$\{(B, b), (R, b)\}$	$\{(B, f), (R, b)\}$		
I	$\{(B, b), (R, b)\}$	0, 0	0, 0	1, -1	1, -1
	$\{(B, b), (R, f)\}$	0, 0	$-\frac{1}{4}, \frac{1}{4}$	$\frac{3}{4}, -\frac{1}{4}$	0, 0
	$\{(B, f), (R, b)\}$	$-\frac{3}{4}, \frac{3}{4}$	0, 0	$-\frac{1}{4}, \frac{1}{4}$	0, 0
	$\{(B, f), (R, f)\}$	-1, 1	-1, 1	-1, 1	-1, 1

payoff to I from this combination of strategies is $(\frac{1}{4})(-1) + (\frac{1}{4})(2) + (\frac{1}{4})(1) + (\frac{1}{4})(-1) = \frac{3}{4}$. Similarly, II 's expected payoff is $(\frac{1}{4})(1) + (\frac{1}{4})(-2) + (\frac{1}{4})(-1) + (\frac{1}{4})(1) = -\frac{1}{4}$. The payoffs for the other cells are calculated in the same way.

Best replies and Nash equilibria

The notion of a player's *best reply* or *best response* is crucial to defining a game's Nash equilibria. Continuing to work with the normal form, suppose that there are M players. Viewing the game from player i 's perspective, the plans of the other players, which are denoted by $s_{-i} = (s_1, s_2, \dots, s_{i-1}, s_{i+1}, \dots, s_M)$, give almost a complete plan for playing the game. It tells how every player other than i will play. Then, a best reply for i to s_{-i} is a strategy that gives i its highest payoff given that the other players are playing according to s_{-i} . If, for example, II 's pure strategy is to stand firm in the game of chicken in Figure A5(a), then I 's best reply is to submit. This strategy leaves I with -1 , whereas standing firm would give -5 . Sometimes a player has more than one best response. If I 's strategy in the simple game of poker in Figure A6 is always to bid, that is, to play $\{(B, b), (R, b)\}$, then II has two best replies. Always bidding or bidding only with a black card, that is, $\{(B, b), (B, b)\}$ or $\{(B, b), (R, f)\}$, will yield II its highest payoff of zero given that I is following the strategy of always bidding. (This can be seen easily by looking across the row associated with I 's strategy of always bidding. In this row, the highest payoff II can attain is zero, and any column or strategy that gives II this payoff is a best reply.) In sum, a player's best reply to a combination of the other players' strategies is a strategy that will maximize this player's payoff given that the other players are following this combination of strategies.

A *Nash equilibrium* of a game is a complete plan for playing the game such that each player's strategy is a best reply to the other players' strategies. That is, the combination $(s_1^*, s_2^*, \dots, s_M^*)$ is a Nash equilibrium if s_i^* is a best reply to s_{-i}^* for every player i . A reason for calling a combination of strategies that has this property an equilibrium is that no player has an incentive to change what it is doing by following some other strategy. Player i has no incentive to deviate from s_i^* given that the other players are following s_{-i}^* , because s_i^* is a best response to s_{-i}^* , and, by definition, a player's best reply to a combination of strategies maximizes its payoff given that the other players follow this combination of strategies. If, however, a combination of strategies, say $(s'_1, s'_2, \dots, s'_M)$, did not satisfy the Nash property that every player's strategy is a best reply to the other players' strategies, then there would be at least one player, say k , such that s'_k would not be a best reply to s'_{-k} . Thus, k could increase its payoff by deviating from s'_k by actually playing a best reply to s'_{-k} . In brief, no player has an incentive

to deviate from its strategy if and only if the strategies form a Nash equilibrium.

The game of chicken in Figure A5(a) has three Nash equilibria. In the first, I stands firm, and II submits. This combination of strategies corresponds to the cell in the upper-right corner. Clearly, I has no incentive to deviate from F by playing S , for I 's payoff to playing S , given that II is playing S , would drop from 1 to 0. Similarly, II has no incentive to deviate from S given that I is playing F , for if it played F , its payoff would fall from -1 to -5 . In the second equilibrium, II stands firm, and I submits. This is the combination at the lower left. As in the previous case, no state has an incentive to deviate from its strategy.

The third equilibrium involves mixed strategies. Suppose that each player will stand firm with probability 0.2 and submit with probability 0.8; then each player's strategy is a best response to the other's, and therefore this combination is a Nash equilibrium. To see that I 's strategy is a best reply to II 's, calculate I 's expected payoff to standing firm: This is I 's payoff if both I and II stand firm times the probability that II will stand firm plus the payoff if I stands firm and II submits times the probability that II will submit. This is $0.2(-5) + 0.8(1) = -0.2$. Similarly, I 's payoff to submitting is $0.2(-1) + 0.8(0) = -0.2$. This shows that if II stands firm with probability 0.2 and submits with probability 0.8, then the payoffs to I of standing firm and of submitting are the same. Thus, I is indifferent to its pure strategies of standing firm or submitting. Indeed, I is indifferent among its mixed strategies as well, for if I stands firm with probability ρ and submits with probability $1 - \rho$, then its expected payoff will be ρ times the expected payoff of standing firm, which is -0.2 , plus $1 - \rho$ times the expected payoff of submitting, which is also -0.2 . This leaves $\rho(-0.2) + (1 - \rho)(-0.2) = -0.2$, regardless of the value of ρ . In sum, I is indifferent among all of its strategies, both pure and mixed. Consequently, all of I 's strategies are best replies to II 's strategy of standing firm with probability 0.2 and submitting with probability 0.8. In particular, I 's strategy of standing firm with probability 0.2 and submitting with probability 0.8 is a best response to II 's strategy.

Just as II 's strategy of standing firm with probability 0.2 left I indifferent among all of its strategies, I 's strategy of standing firm with probability 0.2 leaves II indifferent to all of its strategies. All of II 's strategies are best responses to I 's strategy. Thus, each player's strategy is a best reply to the other's; so the combination of strategies forms a Nash equilibrium.

In general, a finite game, that is, a game that has finite numbers of players and pure strategies, has at least one Nash equilibrium.⁴ But there may

⁴ See Ordeshook (1986, pp. 120–37) and Tirole (1988, pp. 444–5) for a proof of the existence of at least one Nash equilibrium in a finite game.

not be an equilibrium in pure strategies; a Nash equilibrium may exist only in mixed strategies. The matching-pennies game illustrates this. No combination of pure strategies forms a Nash equilibrium. For example, in the combination in which *I* plays *H* and *II* plays *T*, then, given *II*'s strategy of *T*, *I*'s best reply is to deviate from *H* by playing *T*. Although there are no pure-strategy equilibria, there is a mixed-strategy equilibrium in which each player plays *H* with probability $\frac{1}{2}$. If *II* follows this strategy, then *I* will be indifferent between *H* and *T* and all mixed strategies. All of *I*'s strategies are best replies, and, in particular, the strategy of playing *H* with probability $\frac{1}{2}$ is a best response. But if *I* follows this strategy, then *II* is indifferent among all of its strategies. So *II*'s playing *H* with probability $\frac{1}{2}$ is a best reply. Thus, this combination of strategies is a Nash equilibrium.

The mixed strategies illustrate an important fact that is useful in finding the equilibria of the brinkmanship and limited-retaliation models in Chapters 3 through 7. If a player is mixing over two strategies in equilibrium, then both of these strategies must be best replies and consequently provide the same payoff. That is, if a player *i* plays a pure strategy s_i^1 with probability $p > 0$ and another pure strategy s_i^2 with probability $q > 0$, then both s_i^1 and s_i^2 must be best responses, and the utility of playing s_i^1 must equal the utility of playing s_i^2 . If these strategies did not yield the same utility, then one would be preferred to the other. That is, the utility of playing one of the strategies, say s_i^1 , would be greater than the utility of playing s_i^2 . This would mean that the player could increase its payoff by deviating from the mixed strategy in which it plays s_i^1 with probability p and s_i^2 with probability q by choosing a strategy in which it would play s_i^1 with probability $p + q$ and s_i^2 with probability zero. But, by definition, no state can improve its payoff in equilibrium by deviating from its equilibrium strategy. So it must be that s_i^1 and s_i^2 yield the same payoff. Similarly, these strategies must also be best replies, for if they were not, then the player would also be able to increase its payoff by not playing either of them, but playing instead a best reply with probability $p + q$.

The mathematical appeal of mixed strategies is clear. Without them, many games would have no equilibrium. Allowing mixed-strategy equilibria assures that an equilibrium exists. But the empirical meanings and interpretations of mixed strategies and mixed-strategy equilibria are fraught with difficulties.⁵ To illustrate some of these, consider the more general game of chicken in Figure A7, where the numerical payoffs in Figure A5(a) have been replaced by variables. The payoff to standing firm if the other player submits is w , the payoff to submitting if the other player

stands firm is s , the payoff to the compromise outcome that obtains if both players submit is c , and the payoff to the disaster that occurs if both stand firm is d . The game will be one of chicken as long as the payoffs satisfy the following relation: The payoff to prevailing is greater than the payoff to compromising, which is greater than the payoff to submitting, which is better than the payoff to disaster: $w > c > s > d$ for both players *I* and *II*.

Now consider the mixed equilibrium in which *I* stands firm with probability ϕ_I and *II* stands firm with probability ϕ_{II} . To calculate ϕ_I , note that *II*'s expected payoff to standing firm is the payoff to its standing firm and *I*'s standing firm, d_{II} , times the probability that *I* will stand firm, ϕ_I , plus *II*'s payoff if it stands firm and *I* submits, w_{II} , times the probability that *I* will quit, $1 - \phi_I$. This is $d_{II}\phi_I + w_{II}(1 - \phi_I)$. Similarly, *II*'s expected payoff to submitting is $s_{II}\phi_I + c_{II}(1 - \phi_I)$. But now recall that because *II* uses a mixed strategy in equilibrium, *II* must be indifferent between standing firm and submitting. (If it strictly preferred one of these alternatives, then it could improve its payoff by deviating from its mixed strategy to the preferred pure strategy.) *II*'s indifference implies that the expected payoff to standing firm equals the payoff to submitting: $d_{II}\phi_I + w_{II}(1 - \phi_I) = s_{II}\phi_I + c_{II}(1 - \phi_I)$. Solving for the probability that *I* will stand firm gives $\phi_I = (w_{II} - c_{II}) / [(w_{II} - c_{II}) + (s_{II} - d_{II})]$. Similarly, the chances that *II* will stand firm are $\phi_{II} = (w_I - c_I) / [(w_I - c_I) + (s_I - d_I)]$.


The mixed equilibrium has some intuitively appealing properties. One would expect a compromise to be more likely the higher the payoff to compromise, the greater the cost of disaster, and the smaller the payoff to prevailing. The mixed equilibrium conforms to these expectations. The chance of a compromise outcome is the probability that both *I* and *II* will

Figure A7. A more general game of chicken.

		<i>II</i>	
		Stand firm	Submit
<i>I</i>	Stand firm	d_I, d_{II}	w_I, s_{II}
	Submit	s_I, w_{II}	c_I, c_{II}

⁵ For further discussion of this and some attempts to justify mixed equilibria, see Luce and Raiffa (1957, pp. 74–6), Harsanyi (1973), and Harsanyi and Selten (1988, pp. 14–15).

Figure A8. The massive-retaliation game.



submit: $(1 - \phi_I)(1 - \phi_{II})$. This probability increases as the payoff to compromise rises or as the payoffs to disaster and prevailing fall.⁶

But much about this mixed equilibrium is not especially appealing intuitively. Note that the probability that I will stand firm, ϕ_I , does not depend on I 's payoffs, but on II 's. Thus, if I 's payoff to prevailing increases, I 's strategy does not change. Rather, II becomes more likely to stand firm; ϕ_{II} rises as w_I increases. The mathematical reason for this is that in a mixed equilibrium, I 's strategy must keep II indifferent between standing firm and submitting. If, therefore, II 's payoffs do not change, as they do not when only I 's payoff to prevailing rises, then I 's strategy cannot change, for otherwise II would no longer be indifferent. Instead, II 's strategy must change in order to keep I indifferent. A higher payoff to prevailing tends to raise I 's expected payoff to standing firm. This, however, can be offset and I 's indifference restored if II becomes more likely to stand firm, for that will make the prospect of disaster more likely if I stands firm and thus will tend to lower I 's expected payoff to standing firm.⁷

Although the mathematical reasons for these interactions are clear, what, if any, empirical interpretation to attach to them is not so clear, and the interpretations offered in Chapters 3 through 7 must be treated cautiously. One approach to building confidence in any finding is to see if it holds in a wide variety of models. This is very much in keeping with the most important objective of this volume, which is to articulate a general analytic

⁶ $\partial(1 - \phi_I)(1 - \phi_{II})/\partial c > 0$; $\partial(1 - \phi_I)(1 - \phi_{II})/\partial d < 0$; $\partial(1 - \phi_I)(1 - \phi_{II})/\partial w < 0$.

⁷ This argument does not apply to games with more than two players. In those games, a player's mixed strategy may depend on its payoffs.

perspective on nuclear deterrence theory that will help point the way to richer and better models.

Returning to the extensive form, a Nash equilibrium is a combination of behavioral strategies in which each player's behavioral strategy is a best reply to the other players' behavioral strategies. This, however, raises a question. Does it matter whether one analyzes a game in terms of mixed behavioral strategies, in which a player may randomize over the alternatives at each of its information sets, or in terms of mixed strategies, in which a player randomizes over complete plans? If the game is one of *perfect recall*, as the games in this volume are, then these two formulations are equivalent, and the adjective "behavioral" will generally not be used.⁸

Subgame perfection

The game in Figure A8(a) is a simple formulation of the doctrine of massive retaliation when both the United States and the Soviet Union have secure second-strike forces. The Soviet Union begins the game by deciding whether or not to challenge the status quo. If there is no challenge, the status quo continues, and the game ends with payoffs $(0, 0)$. If the Soviet Union challenges the status quo, the United States must decide what to do. It can either carry out a massive nuclear attack or submit by acquiescing to the Soviet challenge. If the United States attacks, the Soviet Union is assumed to retaliate in kind. The game ends in a general nuclear exchange, with payoffs of $(-10, -10)$. If the United States submits, then the United States suffers a loss, and the Soviets gain. The payoffs to this are taken to be $(-8, 8)$.

The normal form of this game is illustrated in Figure A8(b). The game has two Nash equilibria in pure strategies. In the first, the United States plays A , which is a threat to launch a massive nuclear attack if the Soviet Union challenges the status quo, and the Soviet Union accepts the status quo by playing $\sim C$. There is no challenge in this equilibrium. In the second pure-strategy equilibrium, the Soviet Union challenges the status quo by playing C , and the United States acquiesces with S .

⁸ A game is one of perfect recall if no player ever forgets what it previously knew and did. If one thinks of bridge as a two-player game in which each player is playing two hands, then bridge is a game in which there is not perfect recall. When a player is playing one hand, it cannot "remember" its other hand, which it knew when it was bidding that hand. More formally, a game has perfect recall if for any two decision nodes x and y that are in the same information set belonging to a player k , if x' is a decision node preceding x that is in one of k 's information sets, then there must also be a node y' that precedes y and is in the same information set as x' , and the paths leading from x' to x and from y' to y must follow the same alternatives at x' and y' . For a discussion of perfect recall and of the equivalence of these two formulations, see Luce and Raiffa (1957, pp. 159–62) or Selten (1975).

Although both equilibria are Nash (i.e., each state's strategy is a best response to the other's strategy), the first seems implausible as a solution to the game. The American strategy of *A* seems inherently incredible. If, in the tree in Figure A8(a), the United States must actually follow through on its threat by playing *A*, its payoff will be -10 . But if the United States submits, it will receive -8 . Assuming that the United States will act to maximize its payoff whenever it must actually act, then it will play *S* rather than *A*. Accordingly, an equilibrium based on the Soviet Union's believing that the United States will play *A* would seem to be an unreasonable solution for the game.⁹


Much work in game theory has been devoted to refining the notion of an equilibrium by imposing additional restrictions on combinations of strategies beyond the Nash criterion that each strategy be a best reply to the other strategies. These restrictions are intended to exclude unreasonable equilibria like the one just examined from the set of acceptable solutions to the game. One of the simplest restrictions is to demand that a solution be subgame perfect.

Before defining subgame perfection, a subgame must be described. A *subgame* is piece of a game tree that is itself a well-defined game. To find a game's subgames, start with the game's extensive form. Then pick any node in the tree and examine that node and all of the nodes in the tree that come after it. This set of nodes is informationally isolated from the rest of the tree if no information set contains some members of this set of nodes and some nodes in the rest of the tree. If this set of nodes is informationally isolated, then this set of nodes forms a well-defined game beginning at the original node and constitutes a subgame of the original game.

Consider, for example, the American decision node in the massive-retaliation game in Figure A8(a). This decision node and its successors, of which there are none, are informationally isolated. No information set connects the American node with the rest of the tree. A well-defined, albeit very simple, game begins at the American decision node. Accordingly, a subgame begins at this node. The tree in Figure A9 provides another example. A subgame begins at each of *I*'s decision nodes. For the same reasons outlined for the massive-retaliation game, a subgame begins at the two nodes where *I* must choose between *T* and *B*. A subgame begins at *I*'s first decision node because every game is a subgame of itself. This follows,

⁹ It might at first seem that the United States would have an incentive to deviate from its strategy of playing *A* and thus that the combination of strategies (*A*, $\sim C$) could not be a Nash equilibrium. But if the Soviet Union does play $\sim C$, then the American decision node in the tree is never reached. Regardless of what the United States would do if this node were reached, the United States will receive zero because the Soviet Union does not challenge the status quo. Every American strategy is a best reply to the Soviet strategy of not challenging the status quo. Thus, there is no incentive for the United States to deviate from *A* if the Soviet Union plays $\sim C$.

Figure A9. Some examples of subgames.



rather vacuously, from the definition of a subgame, for the first decision node (along with all the nodes that follow it) is informationally isolated from the rest of the tree, because there is no rest of the tree.

But a subgame does not begin at either of *II*'s decision nodes. *II*'s upper node and its successors are informationally linked to the rest of the tree, because one of these nodes, *III*'s upper node, is in an information set containing nodes in the rest of the tree, namely, *III*'s lower node. Thus, a well-defined game does not begin at *II*'s upper node, and so a subgame does not begin there. For similar reasons, a subgame does not begin at *II*'s lower node or at either of *III*'s decision nodes.

Given this description of a subgame, a *subgame perfect equilibrium* can be defined. A combination of strategies forms a subgame perfect equilibrium if the strategies form a Nash equilibrium in every subgame of the original game. In effect, requiring an equilibrium to be subgame perfect means that no player can threaten to play a strategy that is inherently incredible in the sense that this player has an incentive to deviate from this strategy in some subgame. A player cannot threaten to do something in a subgame when doing something else in that subgame would make the player better off. The strategy embodying such a threat would not be Nash in this subgame and so could not be part of a subgame perfect equilibrium. In this way, focusing on subgame perfect equilibria eliminates some unreasonable equilibria.¹⁰

¹⁰ Because every game is a subgame of itself, and a subgame perfect equilibrium is Nash in every subgame, a subgame perfect equilibrium is also a Nash equilibrium. This means that the set of subgame perfect equilibria is a subset of the set of Nash equilibria.

To show that looking for subgame perfect equilibria eliminates the unreasonable equilibrium in the massive-retaliation game, first note that there are two subgames of this game. The first is the game itself, and the second is the subgame beginning at the American decision node. Now consider the strategy $(A, \sim C)$, in which the United States would attack if the Soviet Union challenged the status quo, but the Soviet Union does not dispute the status quo. As shown earlier, this set of strategies is a Nash equilibrium in the original game and therefore is also Nash in the first subgame. But this combination of strategies is not Nash in the subgame beginning at the American node. In this very simple subgame, the United States has an incentive to deviate from A . Playing A will give -10 , and playing S will bring -8 ; the United States' best reply is to submit. Because the combination of strategies $(A, \sim C)$ is not Nash in all subgames, it is not a subgame perfect equilibrium. Thus, looking for subgame perfect equilibria rather than simply Nash equilibria will exclude the unreasonable equilibrium in the massive-retaliation game.

The other equilibrium of the massive-retaliation game, (S, C) is, however, subgame perfect. As demonstrated previously, this combination is Nash in the first subgame of the massive-retaliation game. It is also Nash in the second subgame. In the subgame beginning at the American node, the United States has no incentive to deviate from its strategy of S .

In sum, analyzing a game in terms of subgame perfect equilibria rather than solely in terms of Nash equilibria helps to eliminate some unreasonable Nash equilibria that seem to be based on inherently incredible threats.

Sequential equilibria

Requiring solutions of a game to be subgame perfect excludes some implausible equilibria. But subgame perfection is limited by the fact that many games cannot be cut into very many subgames, because the informational complexity of the games means that few sections of the game tree are informationally isolated from the rest of the tree. In such games, even subgame perfect equilibria may depend on what seem to be inherently incredible threats.

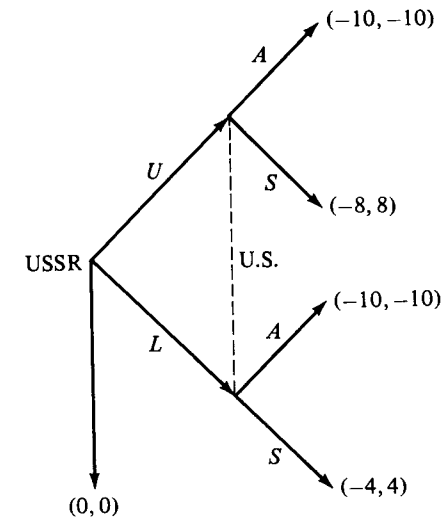
Consider the game in Figure A10. The Soviet Union has three alternatives at the beginning of the game. If it does not challenge the status quo, $\sim C$, the game ends with the status quo payoffs $(0, 0)$. The Soviet Union may also pursue a limited strategy, L , or an unlimited strategy, U . If the Soviet Union pursues a limited strategy and the United States then submits, the payoffs will be $(-4, 4)$. If, however, the Soviet Union is pursuing an unlimited strategy and the United States acquiesces, then the United States

will pay a higher cost. Here the payoffs are $(-8, 8)$. Whether the Soviet Union pursues a limited or unlimited strategy, the United States can launch a massive nuclear attack, A , which will end the game with $(-10, -10)$. Finally, when the United States must decide whether to attack or submit, it does not know whether the Soviet strategy is limited or unlimited. This means, formally, that both of the American decision nodes are in the same information set.

The combination of strategies $(A, \sim C)$ in which the Soviet Union does not challenge the status quo and the United States attacks if there is a challenge is a Nash equilibrium. Given the Soviet strategy of $\sim C$, the American payoff is zero regardless of what it does. Every American strategy is a best response, and, in particular, A is a best reply. Given the American strategy of A , the best the Soviet Union can do is not challenge the status quo: $\sim C$ is the Soviet Union's best reply. Because each player's strategy is a best response to the other's strategy, $(A, \sim C)$ is a Nash equilibrium.

This combination of strategies is also subgame perfect. To see this, note that the game in Figure A10 has only one subgame, which is the game itself. A subgame does not begin at either American decision node, because the part of the tree beginning at either of these nodes is not informationally isolated from the rest of the tree. The United States' information set links the part of the tree beginning at one of the American decision nodes with the rest of the tree. Because $(A, \sim C)$ is Nash in all of the game's subgames,

Figure A10. A game with only one subgame.



which in this case amounts to being Nash only in the game itself, $(A, \sim C)$ is subgame perfect.

Although this combination of strategies is subgame perfect, the equilibrium does not seem reasonable. The American strategy of playing A seems incredible. Just as it seemed implausible that the United States would attack in the massive-retaliation game in Figure A8(a), because if it actually had to act it would always do better by submitting, it also seems unreasonable for the United States to attack in the game in Figure A10. Whether the United States is at its upper or lower node, submitting always offers a higher payoff than attacking. Attacking at the upper node in the information set would bring -10 , and submitting would bring -8 . Attacking at the lower node would also yield -10 , but acquiescing would be even less costly, giving -4 . Accordingly, an equilibrium based on a Soviet assumption that the United States will play A would seem to be an unreasonable solution for the game.

Sequential equilibria may in part be seen as an attempt to exclude equilibria like $(A, \sim C)$ by extending the basic idea underlying subgame perfection.¹¹ Subgame perfection requires that each player behave reasonably in all subgames in the sense that no player can have an incentive to deviate from its equilibrium strategy in any subgame. Clearly, the United States in Figure A10 has an incentive to deviate from its strategy of A if it ever actually has to act. But because a well-defined subgame does not start at this information set, the criterion of acting reasonably in all subgames cannot rule out this American strategy. Suppose, however, one could define a player's payoffs beginning at any information set, not just from a single node at the start of a subgame. Then, just as subgame perfection requires that no player have an incentive to deviate from its strategy in any subgame, one might require that no player have an incentive to deviate from its strategy at any information set given the other players' strategies. This requirement would then rule out an equilibrium like $(A, \sim C)$ in the game in Figure A10, for the United States would always have an incentive to deviate from A . In effect, a sequential equilibrium first specifies a way of calculating a player's payoffs not just within a subgame but starting at any one of its information sets. Then a sequential equilibrium demands that no player have any incentive to deviate from its equilibrium strategy at any of its information sets.

To make this description of a sequential equilibrium meaningful, a way of calculating a player's payoffs starting from any information set must be defined. Suppose a player wanted to calculate the expected payoff of following a specific strategy starting from one of its information sets and

¹¹ See Kreps and Wilson (1982b) and Kreps and Ramey (1987) for a discussion of sequential equilibria.

given the other players' strategies. If the player knew where it was in this information set, then calculating this strategy's expected payoff would be easy. The player could simply trace the path through the tree starting from this node and specified by this player's strategy and the other players' strategies. Consider, for example, the problem confronting I in the simple poker game in Figure A4 if it wants to determine the expected payoff to bidding given that it has drawn a red card and that II 's strategy is to bid if it has a black card and to fold if it has a red card. I knows that it has a red card, but does not know if it is at the upper-right node or the lower-left node in the information set associated with Nature dealing I a red card. If, however, I knows that it is at the upper-right node, that is to say that Nature has actually followed the branch (R, R) , then I can easily calculate the expected payoff of bidding, given II 's strategy. If I bids, II 's strategy is to fold, because II is holding a red card. I 's expected payoff is 1. Similarly, if I knows that it is at the lower-left node [i.e., Nature has dealt (R, B)], then I 's expected payoff to bidding, given II 's strategy (which, if holding a black card, is to bid), is -2 . The problem in calculating the expected payoff of following a particular strategy at a specific information set is that a player does not know where it is in this information set. In the simple poker game, I does not know whether it is at its upper-right node or lower-left node. But suppose that a player has some beliefs about where it is in an information set. That is, a player attaches some probability to being at a specific node given that this player is somewhere in this information set. Then the expected value of following a specific strategy at this information set is the sum over all of the nodes in this information set of the probability of being at any given node times the expected utility of following this strategy starting from this node. I may, for example, believe, after drawing a red card, that the probability that it is at its upper-right node in the simple poker game is $\frac{1}{4}$. Thus, the expected payoff to bidding at this information set is the probability of being at the upper-right node times the expected payoff of bidding at this node, which is 1, plus the probability of being at the lower-left node times the utility of bidding there. This is $(\frac{1}{4})(1) + (\frac{3}{4})(-2) = -\frac{1}{2}$. In sum, once a player's beliefs about where it is in an information set are specified, then this player's expected payoff to following some strategy, given the other players' strategies, can be calculated.

To generalize this way of calculating the expected payoff at a player's information set, let i be some player in an arbitrary game. Player i is assumed to have a *system of beliefs*, which is denoted by μ_i . For each of i 's information sets, μ_i specifies the probability with which i believes that it is at a particular node given that the play of the game has reached the information set containing this node. More formally, μ_i specifies the probability of being at each node conditional on being in the information

set containing this node. In the simple poker game, a system of beliefs for I would define the probability that I would be at the upper-right node and the lower-left node given that I was at the information set associated with its holding a red card. I might, for example, believe, as before, that these probabilities were $\frac{1}{4}$ and $\frac{3}{4}$, respectively. I 's system of beliefs, μ_I , would also have to specify what I would believe should it find itself holding a black card. Recalling that each player is assumed to have a system of beliefs, let μ denote the set of all the players' belief systems. In the simple poker game, $\mu = \{\mu_I, \mu_{II}\}$. Accordingly, μ specifies for each node in the game the probability that the player who owns this node attaches to being at this node given that the play of the game has reached the information set containing this node. Now let (μ, π) be an *assessment* of a game, where μ is a system of beliefs and π is a combination of the players' strategies that provides a complete plan for playing the game. An assessment contains enough information to permit the calculation of a player's expected utility to following a particular strategy at any one of its information sets. With π , one can calculate any player's expected payoff to following this strategy starting from a specific node in this information set. Then, with μ specifying the relative likelihood of being at a particular node in this information set, one can calculate the expected payoff to following this strategy at this information set, as was done earlier in the poker-game example.

A *sequential equilibrium* can now be defined as a special type of assessment. More specifically, an assessment (μ, π) is a sequential equilibrium if it satisfies two conditions. The first is that the assessment must be *sequentially rational*. This means that no player has an incentive to deviate from its strategy at any one of its information sets given its beliefs and the other players' strategies. This is merely the extension of the basic idea underlying subgame perfection.

To clarify what it means to be sequentially rational, consider the following assessment. I 's strategy is to bid if dealt a black card and to fold with a red card. II 's strategy is always to bid: $\pi = (\pi_I, \pi_{II}) = (\{(B, b), (R, f)\}, \{(B, b), (R, b)\})$. Suppose further that I believes that if it is holding a red card, the chance that II 's card is black is $\frac{1}{2}$, and therefore the probability that II 's card is red is also $\frac{1}{2}$. Or, equivalently, given that play has reached the information set belonging to I at which I holds a red card, then the probability of actually being at the lower-left node in this information set is $\frac{1}{2}$. Similarly, I also believes that if it has been dealt a black card, then the probability that II has been dealt a red card is $\frac{1}{2}$, and the chance that it has a black card is also $\frac{1}{2}$. II 's beliefs are simpler. Regardless of what its card is, II is certain that I has a red card. Momentarily setting aside the question of whether or not these beliefs are reasonable, the assessment composed of this system of beliefs and strategies is sequentially rational.

To be sequentially rational, no player can have any incentive to deviate from its strategy given its beliefs and the other players' strategies. II clearly has no reason to change its strategy given its beliefs. Believing that I is certain to be holding a red card, bidding brings 2 if II 's card is black, and 0 if II 's card is red. Folding always brings -1 . Given II 's beliefs, bidding is its best reply. I also has no incentive to alter its strategy given its beliefs and II 's strategy. Given that II will always bid, I 's payoff to bidding if it has a black card is 0 if II actually has a black card, and 2 if II 's card is red. I believes that the probability that II 's card is black is $\frac{1}{2}$; so I 's expected payoff to bidding is $(0)(\frac{1}{2}) + (2)(\frac{1}{2}) = 1$. If, however, I deviates by folding with a black card, its payoff will be -1 . If, instead, I tries a mixed strategy of bidding with probability p , then this strategy's payoff is the probability of bidding times the expected payoff to bidding plus the probability of not bidding times the payoff to that. So a mixed strategy yields $p(1) + (1-p)(-1) = 2p - 1$, which is also less than or equal to 1. Thus, I cannot improve its payoff by deviating; bidding with a black card is I 's best reply given its beliefs. A similar argument shows that folding with a red card is I 's best response given its beliefs and II 's strategy. No player has any incentive to deviate from its strategy given its beliefs and the other player's strategy; so this assessment is sequentially rational.

Sequential rationality is one of two conditions an assessment must satisfy in order to be a sequential equilibrium. The second condition has to do with the system of beliefs. Just as some Nash equilibria were excluded because the strategies seemed unreasonable, some belief systems seem unreasonable and will be excluded. Indeed, although the assessment just described is sequentially rational, the beliefs underlying it do not seem sensible. When II bids, it is, according to its system of beliefs, certain that I 's card is red. But II will bid only if I has already bid, and I , according to its strategy, will bid only if it has a black card. Given I 's strategy, II should believe that I is holding a black card if and when II has to decide whether or not to bid. II 's beliefs are incompatible with I 's strategy.

The second condition an assessment must satisfy if it is to be a sequential equilibrium is that the belief system must be "reasonable" in the sense that it is *consistent*.¹² Requiring beliefs to be consistent entails a number of subtleties and difficulties.¹³ Fortunately, the games analyzed in the

¹² The questions what constitute "reasonable" beliefs and, more generally, how to "refine" Nash equilibria in order to eliminate the unreasonable ones have motivated an immense amount of recent work in game theory. For further discussion of this, see Kreps and Wilson (1982b), Rubinstein (1985), Grossman and Perry (1986), Banks and Sobel (1987), Kreps and Ramey (1987), and Cho (1987).

¹³ See Kreps and Wilson (1982b) and Kreps and Ramey (1987) for the formal definition of a consistent assessment and some of its subtleties.

preceding chapters are sufficiently simple that these difficulties and subtleties do not arise. The only important consistency criterion for the models examined in the preceding chapters is that the system of beliefs satisfy Bayes' rule where this rule can be applied. An assessment like this that is sequentially rational and satisfies Bayes' rule where this rule applies is a *perfect Bayesian equilibrium*.¹⁴ Bayes' rule is a means of revising a prior probability in light of some new information or evidence. In the present context, Bayes' rule provides a way of updating a prior probability of reaching a given decision node in light of play having actually reached the information set containing this node. It provides a way, for example, for *I*, after being dealt a red card, to revise the belief it held before the deal that red cards would be dealt to both it and *II*.

Bayesian updating of beliefs is crucial to understanding the dynamics of the models analyzed in this volume. But before discussing Bayesian updating in a game-theoretic context where strategic interactions must be taken into account, it will be useful to discuss Bayesian updating in a simpler context in which there is only one player and no strategic interaction. Suppose that an urn can be filled with either of two possible mixtures. The urn may contain seventy-five green marbles and twenty-five blue ones, or it may hold twenty-five green marbles and seventy-five blue ones. The player believes that the two mixtures are equally likely. (This probability might be a subjective estimate; it could be based on a statistical analysis of some previously obtained data, or, if guessing the contents of this urn was a rather dull parlor game, then this probability might be due to the way that the mixture was chosen, say by flipping a coin.) Now the player is allowed to draw two marbles. Both are green. Given this new evidence, how should the player update the probability that the mixture is 75 percent green? Bayes' rule provides a means of doing this.

The key to Bayes' rule is to observe that there are two ways of thinking about the probability that two events, say *X* and *Y*, will happen. Let $P(X \cap Y)$ denote the probability that both *X* and *Y* will occur. In the urn example, *X* is the event "two green marbles are drawn," and *Y* is the event "the mixture is 75 percent green." One way to think about the probability that both *X* and *Y* will happen is that this is the same as the probability that *X* will happen, given that *Y* will occur, times the probability that *Y* will happen. The probability that *X* will happen given that *Y* will occur is the *conditional probability of X given Y* and is denoted by $P(X|Y)$. In the example, $P(X|Y)$ is the probability of drawing two green marbles given that the mixture is 75 percent green. This is the probability that the first draw

¹⁴ This is the weakest notion of a perfect Bayesian equilibrium. Stronger ones are obtained by making assumptions about what "reasonable" beliefs are where Bayes' rule cannot be applied.

will be green, which is $\frac{75}{100}$, times the probability that the second marble will be green, which, because there are only ninety-nine marbles left and seventy-four are green, is $\frac{74}{99}$. The probability of drawing two greens is therefore $(\frac{75}{100})(\frac{74}{99}) = 0.561$. Letting $P(Y)$ be the initial or *prior probability* of *Y*, which in this example is the initial probability of a mostly green mixture or $\frac{1}{2}$, then the probability of both *X* and *Y* is equal to the chance of *X* happening, given *Y*, times the probability of *Y* occurring, or $P(X \cap Y) = P(X|Y)P(Y) = (0.561)(\frac{1}{2}) = 0.280$.

But there is another way to think about the chances that both *X* and *Y* will happen. This is also the probability that *Y* will occur, given *X*, times the probability that *X* will happen, or $P(X \cap Y) = P(Y|X)P(X)$. The conditional probability $P(Y|X)$ is, in the example, the probability that the mixture is 75 percent green given that both the drawn marbles are green. This, moreover, is the updated probability that the player is trying to calculate.

To find an expression for this updated probability, bring together the two ways of thinking about the chances that both *X* and *Y* will occur, to obtain $P(Y|X)P(X) = P(X \cap Y) = P(X|Y)P(Y)$. Solving this for the updated probability that the player is trying to calculate, $P(Y|X)$, gives Bayes' rule for updating probabilities: $P(Y|X) = P(X \cap Y)/P(X) = P(X|Y)P(Y)/P(X)$. That is, the probability of *Y*, given *X*, is the probability of *X* and *Y* divided by the prior probability of *X*. Or, in the urn example, the probability of a 75 percent green mixture, given that two greens have been drawn, is the probability of a mostly green mixture and a draw of two greens divided by the prior probability of drawing two greens. These probabilities are readily calculated. The former, as calculated earlier, is the probability of two greens, given a mostly green mixture, times the prior probability of a mostly green mixture, or $(0.561)(\frac{1}{2}) = 0.280$. The prior probability of drawing two greens, $P(X)$, is the probability of two greens, given a mostly green mixture, times the probability of a mostly green mixture plus the probability of drawing two greens from a mostly blue mixture times the probability of a mostly blue mixture. This is $(\frac{75}{100})(\frac{74}{99})(\frac{1}{2}) + (\frac{25}{100})(\frac{26}{99})(\frac{1}{2}) = 0.311$. Thus, the Bayesian update of the chance that the mixture is mostly green after two green marbles have been drawn is $0.280/0.311 = 0.902$. After drawing two green marbles, the prior probability that the mixture was mostly green, which was $\frac{1}{2}$, has been updated to 0.902.

Returning now to a game-theoretic context, consider *I*'s beliefs in the sequentially rational assessment described earlier for the simple poker game. They are consistent with Bayes' rule, as they must be in a sequential equilibrium or in a perfect Bayesian equilibrium. The prior probability of being at any one of *I*'s decision nodes is $\frac{1}{4}$. That is, before the deal, *I*'s estimate or prior probability of being at a specific decision node, say the node associated with Nature's dealing (*R*, *R*), is $\frac{1}{4}$. But after the deal, *I* knows

that it is holding a red card and can then revise its beliefs to incorporate this new information. According to Bayes' rule, the probability of being at (R, R) , given that play has reached I 's information set associated with I 's holding a red card, is the prior probability of being at (R, R) , which is $\frac{1}{4}$, divided by the probability that the play of the game will reach this information set. This latter probability is simply the sum of the probabilities of reaching all of the individual nodes in this information set or, in this case, the probability of reaching (R, R) plus the probability of reaching (R, B) , which is $\frac{1}{2}$. Bayes' rule assigns a probability of $(\frac{1}{4})/(\frac{1}{4} + \frac{1}{4}) = \frac{1}{2}$ to being at (R, R) , given that I knows it is holding a red card. This is precisely what I 's belief system in the sequentially rational assessment says that I believes about II 's card, given that I is holding a red card. I 's beliefs are consistent with Bayes' rule.

Now consider II 's beliefs. When bidding, II believes that I is certain to be holding a red card. As noted earlier, this belief seems unreasonable, given I 's strategy, because II will have an opportunity to bid only if I bids, and I will bid only if its card is black. Indeed, the only thing that it seems reasonable for II to believe about I 's card, given I 's strategy, is that I 's card is black. Requiring beliefs to be consistent with Bayes' rule simply formalizes this reasoning, and this shows that II 's beliefs do not conform to Bayes' rule. The sequentially rational assessment described earlier is therefore not a sequential equilibrium.

To see that II 's beliefs are incompatible with Bayes' rule, suppose that II holds a black card, and I bids. II , therefore, is somewhere in its lower information set in Figure A4. But where does II believe it is? What, for example, is the probability that it is at the upper-left node? Or, equivalently, what is the probability that Nature has dealt (B, B) and I has bid? The first step in calculating this probability is to find the prior probability of reaching this node (i.e., the probability of reaching this node as calculated before the game begins). This is the probability that Nature will deal (B, B) times the probability that I will bid with this deal. Nature will deal (B, B) with probability $\frac{1}{4}$, and I , according to its strategy, will always bid when dealt a black card. The prior probability of reaching II 's upper-left decision node in its lower information set is $\frac{1}{4}$. Similarly, the probability of reaching the lower-right node in this information set is the probability that Nature will deal (R, B) and that I will bid. This is $(\frac{1}{4})(0) = 0$. The updated probability of being at II 's upper-left node, given that I has actually bid, or, in other words, that play has actually reached the information set containing this decision node, is then obtained by dividing the prior probability of reaching this node by the probability of reaching this information set. This latter probability is $\frac{1}{4} + 0$; so the updated probability is $(\frac{1}{4})/(\frac{1}{4}) = 1$. That is, II , according to Bayes' rule, is certain that it is at its upper-left decision node.

Similarly, II believes that the probability that I is holding a red card when II is actually bidding is $0/(\frac{1}{4} + 0) = 0$, not 1 as in the sequentially rational assessment. Beliefs in this assessment are not in accord with Bayes' rule, and this means that the assessment cannot be a sequential equilibrium.

To state the requirement that beliefs satisfy Bayes' rule somewhat more generally, let y be some decision node, and let h be the information set containing y . Then, for any assessment (μ, π) , the probability of reaching y can be calculated. Let $P(y|(\mu, \pi))$ denote this probability. Similarly, the probability of reaching h can be calculated. It is $P(h|(\mu, \pi)) = \sum_{x \in h} P(x|(\mu, \pi))$, where x is a node in h , and the summation is taken over all of the nodes in h . Then, if $P(h|(\mu, \pi)) > 0$, Bayes' rule says that the probability of being at y , given that play has actually reached h , is $P(y|(\mu, \pi))/P(h|(\mu, \pi))$.

Clearly, Bayes' rule cannot be applied if the probability of reaching an information set is zero [i.e., if $P(h|(\mu, \pi)) = 0$], for trying to use the rule in this case would entail dividing by zero. However, as long as $P(h|(\mu, \pi)) > 0$, Bayes' rule can be used, and the only consistency criterion required of beliefs in the models in this volume is that beliefs satisfy Bayes' rule at information sets where this rule can be applied.¹⁵

Games of incomplete information

The final issue to be discussed is the problem of incomplete information.¹⁶ Players in a situation may have incomplete information about the other

¹⁵ What distinguishes a sequential equilibrium from the weakest notion of a perfect Bayesian equilibrium, which is the one employed here, is that a sequential equilibrium places weak consistency restrictions on beliefs at information sets that are reached with probability zero. To describe a consistent assessment and to specify more formally what conditions consistent beliefs must satisfy at information sets that are reached with zero probability, let π^1 be a completely mixed set of strategies for playing the game. A set of strategies is completely mixed if each participant plays every alternative at each of its information sets with a positive probability. That is, no alternative is played with zero probability in π^1 . Because every alternative is played with positive probability, every information set h is reached with positive probability. Accordingly, Bayes' rule can be applied at every information set in the game. Let $\mu^1(y)$ be the probability of being at y given that play has reached the information set containing y , which will be denoted by $h(y)$. Then, by Bayes' rule, $\mu^1(y) = P(y|\pi^1)/P(h(y)|\pi^1)$. In brief, Bayes' rule can always be used to define a system of beliefs μ^1 when π^1 is completely mixed. An assessment (μ, π) is consistent if and only if there exists a sequence of completely mixed assessments that converges to (μ, π) . Symbolically, there must exist a sequence of $\{(\mu^i, \pi^i)\}_{i=1}^{\infty}$, where the π^i are completely mixed and are such that $\lim_{i \rightarrow \infty} (\mu^i, \pi^i) = (\mu, \pi)$. For a detailed discussion of consistency and some of the subtleties associated with it, see Kreps and Wilson (1982b) and Kreps and Ramey (1987).

¹⁶ Harsanyi (1967–8) originated this approach.

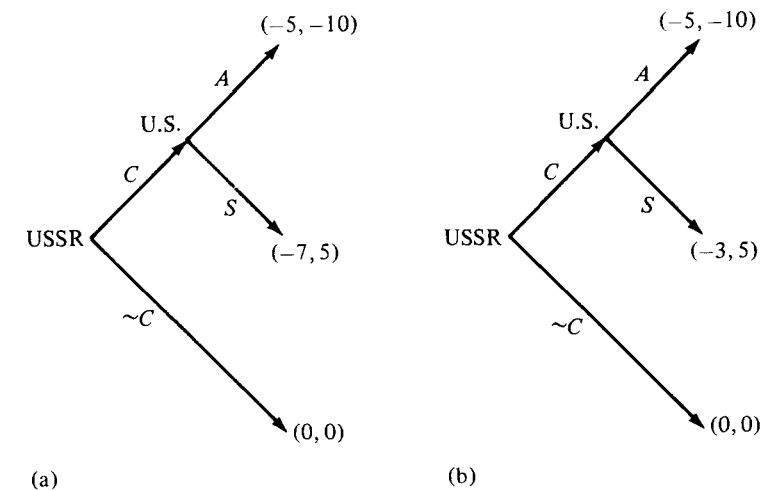
players. A player may be uncertain of the other players' payoffs or of the set of alternatives from which the other players can choose. In crises, for example, states often are said to be unsure of the resolve of their adversaries. That is, a state lacks complete information about its adversary's willingness to run risks or about what the adversary sees as being at stake in the crisis. Games of incomplete information are used to model situations in which players are uncertain about some aspects of the situations confronting them. An important feature about these games is that players can try to learn about the other players by observing what they do. Of course, an adversary, understanding this, also may have an incentive to try to misrepresent its type, to try, for example, to appear to be more resolute than it actually is. Games of incomplete information are used to study these competing influences and their effects on the players' strategies.

An example may be the best way to illustrate how games of incomplete information are set up and analyzed. The example is a variant of the simple model of massive retaliation used earlier in the discussion of subgame perfection. In this variant, the Soviet Union is uncertain of the cost to the United States of acquiescing to a Soviet challenge. Suppose, that is, that when the United States is relatively invulnerable to Soviet retaliation, the United States attempts to prevent a Soviet challenge by threatening to retaliate massively to a Soviet provocation. In the game, the status quo payoffs are $(0, 0)$, and the respective payoffs to the United States and the Soviet Union will be $(-5, -10)$ if there is a Soviet challenge and a massive American nuclear attack in response. (The payoff of -5 reflects an assumed relative American invulnerability.) The Soviet Union will also receive 5 if the United States acquiesces to a Soviet challenge. The Soviet Union, however, lacks complete information about the United States. In particular, the Soviet Union is unsure if the United States attaches a high value to what is at stake, so that submission will bring a large loss of -7 , or if the United States puts a low value on what is at stake, so that submission will bring only a small loss of -3 .

Figure A11(a) shows the tree and payoffs if the cost to giving in is high, and Figure A11(b) depicts the tree and payoffs if the cost is low. The problem would be easy to analyze if the Soviet Union were sure of the American payoff to acquiescing. If the cost to giving in were known to be high, the Soviet Union would be in the game in Figure A11(a), where the unique subgame perfect equilibrium is for the Soviet Union not to challenge and for the United States to attack if challenged. Similarly, if the cost of American acquiescence were known to be low, the game in Figure A11(b) would be the relevant one. Here the unique subgame perfect equilibrium has the Soviet Union challenging the status quo and the United States submitting.

The difficulty is, of course, that the Soviet Union is uncertain whether the cost to the United States of submitting is high or low. To model this lack of complete information, the two games in Figure A11 are combined into a single, larger game. Suppose that the Soviet Union believes that the probability that the United States attaches a high cost to submitting is ρ , and the probability of a low cost is $1 - \rho$. Then the games in Figure A11 may be combined to form the game in Figure A12. This game begins with Nature making a random move. This is the modeling device used to create the Soviet Union's uncertainty about the American cost of submission. If Nature takes the upper branch, which it will do with probability ρ , then the rest of the tree beginning at the Soviet decision node is the same as the tree in Figure A11(a). (The prime on "U.S." indicates that along this path through the tree, the United States attaches a high cost to submitting and will play accordingly.) Thus, if the Soviet Union were certain that it was at the upper node in its information set in the game in Figure A12, this game would be played in exactly the same way as the game in Figure A11(a). Similarly, if Nature takes the lower branch, which it will do with probability $1 - \rho$, then the rest of the game starting from the Soviet Union's lower decision node corresponds to the tree in Figure A11(b). If the Soviet Union were certain that it was at this lower node, then the game in Figure A12 would be played just like the game in Figure A11(b). The Soviet Union, however, does not know if it is at its upper or lower node, for they are in the same information set. Rather, the Soviet Union forms beliefs about where it is in its information set. Following Bayes' rule, the Soviet Union believes

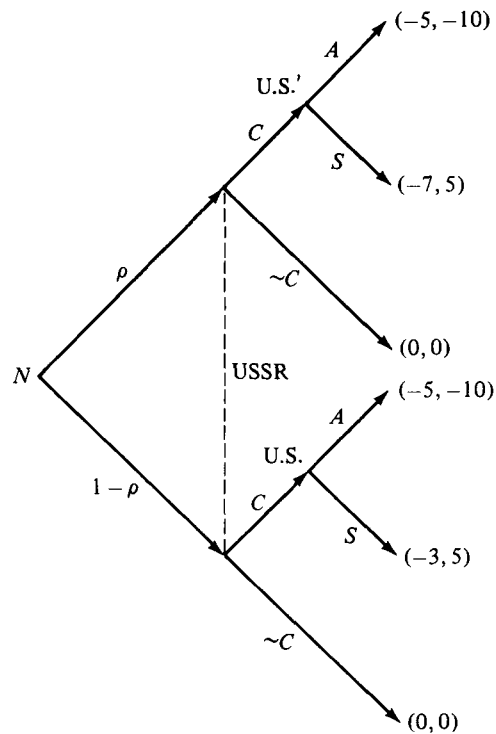
Figure A11. Massive retaliation with high and low stakes.



that it is at its upper node with probability ρ and at its lower node with probability $1 - \rho$. In effect, the Soviet Union begins the game believing that the probability that the United States attaches a high cost to submitting is ρ and that the probability that the United States attaches a low cost is $1 - \rho$. In this way, the larger game in Figure A12 models the Soviet Union's lack of complete information and beliefs about the American payoffs. This game can then be solved for its sequential equilibria, and the equilibrium strategies in this larger game will incorporate the Soviet Union's uncertainty about the American payoffs.

A second example of an incomplete-information game will illustrate the interaction between beliefs and strategies. In this game, which is depicted in Figure A13(a), a potential challenger, C , begins by deciding whether or not to challenge the status quo. If it decides not to mount a challenge, the game ends with continuation of the status quo. If the potential challenger disputes the status quo, the defender, D , can either resist, R , or submit, S . If the defender submits, the game ends. If it resists, then the challenger must decide whether to attack, A , or back down, S .


Figure A12. Massive retaliation with incomplete information.



The status quo payoffs are $(0, 0)$, where the first element of this pair is the challenger's payoff. If the defender submits, the challenger receives 10, and the defender loses 10. If the defender resists and the challenger backs down, the challenger loses 10, and the defender gains 10. If the challenger attacks, the defender's payoff to the ensuing war is -15 . The defender is, however, uncertain of the challenger's payoff to fighting. There are two possibilities. (There could, of course, be more possibilities, but that would make the resulting game difficult to analyze.) The challenger's payoff to attacking may be sufficiently low, say -15 , that it will prefer backing down to attacking if D resists the challenge. These are the payoffs in Figure A13(a). Or the challenger's payoff to fighting may be high enough, say -5 , that it will rather attack than submit if resisted. Figure A13(b) shows these payoffs, where C' denotes the more determined challenger.

As in the massive-retaliation example, the situation would be easy to analyze if the defender were certain of the challenger's payoffs. If, as in Figure A13(a), the challenger's payoff to attacking is so low that it will

Figure A13. Escalation with different payoffs to fighting.




prefer backing down to attacking, then the defender should resist, for the challenger will then submit. Indeed, foreseeing that it will eventually back down, the potential challenger will not even dispute the status quo. The unique subgame perfect equilibrium of the game in Figure A13(a) has the potential challenger accepting the status quo, the defender resisting if challenged, and the challenger backing down. (Remember that an equilibrium describes what will be done at every information set even if in equilibrium some of these information sets are not reached.)

If the defender is certain that the challenger prefers attacking to backing down, then resistance will bring -5 , whereas submitting will cost only 10 . In this case, D will not resist, and the potential challenger will actually challenge the status quo. The unique subgame perfect equilibrium for the game in Figure A13(b) is for the challenger to mount a challenge, the defender to submit, and the challenger to fight should the defender resist.

But the defender is uncertain of the challenger's payoffs. Suppose the prior probability of facing a challenger that prefers fighting is ρ , and the probability of confronting a challenger that would rather quit is $1 - \rho$. The game in Figure A14 represents this situation. Once again, incomplete information is modeled by having Nature begin the game with a random move that leaves D uncertain about the type of its adversary.

Note, however, that what the defender believes about the challenger

Figure A14. Escalation with incomplete information.



depends both on the defender's prior belief and on what the challenger does. This was not an issue in the previous example of an incomplete-information game, because the Soviet Union moved before the United States. Thus, the Soviet Union, which was the state that lacked complete information about its adversary in that example, could not update its beliefs about the United States' payoffs based on what the United States had actually done. This new information was not yet available. In the current example, however, the uncertain state, D , decides what to do after the other state has moved. Accordingly, the defender can update its prior belief about the challenger's willingness to fight in light of the challenger's decision whether or not to challenge the status quo.

To illustrate the interdependence between the challenger's strategy and the defender's updated beliefs, suppose initially that both C and C' are certain to escalate. Intuitively, if C and C' will behave identically, there is nothing to be learned from seeing what the challenger actually does. The updated probability of facing a particular type of challenger will not differ from the prior probability of facing that type of challenger. That is what Bayes' rule shows. The probability of facing the determined challenger C' if D has actually been challenged, according to Bayes' rule, is the prior probability of reaching the lower decision node in D 's information set divided by the probability that play will actually reach this information set. The prior probability of reaching the lower decision node is the prior probability of facing the more determined challenger C' , which is the probability that Nature will follow the lower branch times the probability that C' will mount a challenge. Given that the more determined challenger's strategy is always to mount a challenge, the prior probability of reaching D 's lower node is $\rho \cdot 1 = \rho$. The probability of actually reaching D 's information set, that is, the probability that the potential challenger will really challenge the status quo, is the probability that Nature will take the upper branch, which is $1 - \rho$, times the probability that C will challenge the status quo plus the probability that Nature will follow the lower branch, which is ρ , times the probability that C' will challenge the status quo. This is $\rho \cdot 1 + (1 - \rho) \cdot 1 = 1$. So the probability of facing C' , given the potential challengers' strategies and the fact that the status quo has actually been challenged, is ρ . If both types of challengers will behave in the same way, the challenger's actual behavior reveals nothing about it, and the Bayesian update of the probability of facing a specific type of challenger is unchanged from the prior probability.

But suppose that the two types of challengers will behave differently. Then, observing what has actually happened may say something about the type of the challenger. To illustrate this, assume that the determined challenger still will be certain to dispute the status quo, but the probability

that the less determined challenger will dispute the status quo is 0.1. If, given these strategies, the defender is challenged, it would seem that the chance that the challenger is more determined rather than less is quite high. The updated probability of facing C' rather than C is high. Bayes' rule again formalizes this. The updated probability of facing C' if there has been a challenge is the prior probability of facing C' , which is still ρ , divided by the probability of there being a challenge or, equivalently, of play actually reaching D 's information set. This latter probability is, as before, the prior probability of facing C' times the probability that C' will dispute the status quo plus the prior probability of facing C times the probability that it will challenge the status quo. The updated probability of facing C' , given the potential challengers' strategies and the fact that there has been a challenge, is $\rho/[\rho \cdot 1 + (1 - \rho)(0.1)]$, which is much greater than ρ . (If, for example, $\rho = 0.25$, then the updated probability is $(0.25)/[(0.25) + (1 - 0.25)(0.1)] = 0.77$.)

In this case, the defender has used what has actually happened in the game to revise its beliefs about the type of its adversary. This is a common feature of games of incomplete information. Of course, the challenger realizes that the defender is trying to ascertain the challenger's type by watching what it does. This may create an incentive for the challenger to behave differently than it otherwise would in order to misrepresent its type. C , for example, may want to try to convince the defender that it is facing C' and thus should not resist a challenge. These are some of the issues that games of incomplete information and their sequential equilibria help to illuminate.

In both examples of incomplete-information games there was one-sided incomplete information. Only the Soviet Union was uncertain about some aspects of the United States; the United States was completely certain of the relevant aspects of the Soviet Union. Similarly, only the defender was uncertain of some aspects of the situation in the second example. Nevertheless, the same approach to modeling incomplete information may be extended to the case in which every player is uncertain about some aspects of the other players. Incomplete information can, in general, be modeled by creating a game in which Nature will behave probabilistically, so that each player will begin this game with beliefs that reflect its uncertainty or lack of complete information about the other players. The equilibrium strategies in this often very large and complicated game will then reflect the players' incomplete information and the players' attempts to resolve and exploit this uncertainty.