Chapter One: The Concepts of Game Theory*

1.1 Introduction

Game Theory is the science of strategy. It is a mathematical construct with a clear set of concepts and assumptions and their resulting theorems. And, just like any formal theory, it is subject to the empirical validation of its application to the real world. Issues of strategy arise in economics, business, politics and most of what is the domain of the social sciences. They also arise in biology, although the strategic issues in that discipline are expressed in terms of survival of the fittest rather than in terms of the calculated attempt to prevail that is usually understood to be the purpose of strategy.

Formulating a rational strategy can always be rephrased as an optimization problem. Prevailing or surviving often involves overcoming risks, threats, or the possibly adverse effects of other people's own self-interested behavior. What distinguishes Game Theory from the more general mathematics of optimization is that the individuals involved differ on what needs to be optimized and that the decisions of each affect the objectives of all. Sometime, the individuals involved have convergent interests and optimization can focus on pooling resources or exploiting individual talents for the benefit of the group. At other times, individuals have opposite interests and maximizing one's chances to survive, or to prevail requires minimizing the others' such chances. Sometime, both kinds of interests arise at once. A typical example is two individuals involved in a bargaining situation: they have a common interest in reaching a deal, but each side wants a larger share of what is to be divided. Without agreement, no one gets any deal. But in any agreement, the larger one side's share the smaller the other side's is, and the bigger its loss.

Formalizing issues of strategy always involves answering four questions:

1. Who are the actors?
2. What can they do?
3. What do they want?
4. What do they know?

Game Theory rests on the assumption of rationality: each individual will do the best they can to achieve what they want given what they know. This is not just the use of reason: it is the use of reason with limited means and information, and with a purpose that might not appeal to everyone. In the game-theoretic brand of rationality decision-makers can be thoroughly insane and uninformed. But as long as they dutifully pursue their insane goals given the little they know, they pass the game theorists' test of rationality.

As a formal discipline, Game Theory arguably begins with Antoine Augustin Cournot in the 19th century. Cournot first described the so-called "duopoly" as a market game played by the owners of two water wells.[1] Each well-owner benefits from pumping more water out and selling it, as long as the overall production of water is not too high. Beyond a certain point, depending on current demand, pumping more water out depresses the price so much that it becomes counter-productive. Cournot studied how the two well-owners should behave in order to maximize their income in this competitive environment.

Game theory was named by John von Neumann who first proved that any two-person zero-sum game has an equilibrium point, meaning a pair of strategies, one for each side, such that neither side can improve its own lot by deviating unilaterally from its specified strategy in the pair.[2] A zero-sum game is one where anything that one side wins is the other side's loss. Von Neumann went on to publish an influential textbook with Oscar Morgenstern.[3]

Modern Game Theory perhaps begins with John Nash who, in 1950, expanded von Neumann's result to non-zero sum games involving any number of actors. Nash's contribution was both conceptual and technical: he defined the concept of what is now known as the "Nash equilibrium" and proved its existence in all games that can be formulated in "strategic terms" only. The Nash equilibrium is a set of strategies, one for each actor of the game, such that no one can improve their outcome by deviating unilaterally. This concept remained the centerpiece of Game Theory well into the 1970's.

A very influential text by Duncan Luce and Howard Raiffa appeared in 1957.[4] In that book, the authors make a clear distinction between the so-called "extensive form" and the "normal form" of a game. The extensive form describes precisely what each actor can do and know at any point in the unfolding of a game. The normal form only summarizes their choices as "upfront" strategies which, used against the other actors strategies, yield various outcomes. Nash used the normal form in his result. But it turns out that the extensive form can hold critical information that is not well addressed with the normal form and the Nash equilibrium concept.[5] This was led to two important advances in the 1970's.

Reinhart Selten pointed out that the Nash equilibrium can stipulate decisions that are obviously absurd if the point where such decisions could be implemented is never reached by

---

[1] *Recherches sur les Principes Mathématiques de la Théorie des Richesses*, 1838.

[2] "Zur Theorie der Gesellschaftsspiele," *Mathematische Annalen*, 100(1), 1928.

[3] *Theory of Games and Economic behavior*, with Oskar Morgenstern, 1944.

[4] *Games and Decisions*.

[5] The normal form erases issues of timing and information that are explicit in the extensive form.

expected play. The reason for that failure is that, since such decision points are not expected to ever be reached, the anticipated choices they involve are irrelevant in the very computation of expected outcomes. If such choices are interpreted as threats or pledges they may not be credible since they might not be optimally implemented if tested. Selten proposed a solution to this issue in the form of a refinement of the Nash equilibrium concept that he called the subgame perfect equilibrium (SPE). In essence, the SPE remains a Nash equilibrium in all parts of the game (the so-called subgames) whether these are expected to be reached or not.[6]

John Harsanyi addressed another important concern: at their turns of play, the actors may have limited information about what state of the game they are currently in. One important issue is that these actors might not be entirely sure of what the other actors' priorities or capabilities really are. However, the observation of what the others did previously can hold valuable information on who they really are and what can therefore be expected of them in the future. Harsanyi introduced the use of Bayesian updating, a standard technique from Probability Theory, in the calculation of game equilibrium. It formalizes the fact that if a certain "type" of opponent would be more likely than others to make a certain move then the very realization of that move makes it more likely that one actually faces that type of opponent.

The combination of Nash's, Selten's and Harsanyi's ideas led to a proliferation of solution concepts. The most influential modern concept is the perfect Bayesian equilibrium (PBE) that combines the ideas of both Selten and Harsanyi. But the issues addressed by the various concepts become even more complex when the given game is repeated in time and the actors become interested in the future consequences of their past and present decisions. Formalizing such repeated games involves answering two further questions about the actors:

1. How do they remember the past?
2. How do they appraise the future?

In practice, actors interpret prior history in ways that define a number of "states of the game." Their choices at each turn then define possibly probabilistic transitions from state to state, giving these repeated games a flavor of Markov chains. So-called Markov strategies then define the actors' responses to each state of the game and this leads to the modern concept of Markov perfect equilibrium (MPE). As for the future, it is usually appraised in discounted fashion: the further away in the future the less it matters in present decisions.

These lecture notes will discuss all these successive advances as well as many of their applications. Our focus is on applications and they will rely heavily on the *GamePlan* software as a modeling tool. Some theoretical developments as well as advanced topics inaccessible with *GamePlan* will also appear in later chapters.

[6] Technically, a subgame must be a whole game all by itself.

1.2 The Objects of Game Theory

The actors of a game are formally called the "players." To be a player in a game, you need to have influence over its unfolding and an interest in its possible outcomes. Players are assumed completely independent of each other so that they cannot be bound by the will of others or by enforceable contracts in their own decisions. They are also assumed rational in the strict sense that they aim to fulfill their own interests. But this does not preclude cooperating with some or all of the other players, or even reaching some agreement on how to coordinate actions in profitable ways. However, there is no enforcement mechanism other than the threat of future retaliation for the failure to respect an agreement.

In many games, there is also a special actor, called Chance or Nature, who has no preferences over the outcomes but may act with specified probabilities in certain circumstances. Technically, Chance is therefore not a true player.

As a game proceeds, various states of the world it describes can be reached. The most basic state of a game is called a "node" at which only one player, or Chance, has options to choose from. What the players can choose are called "moves" and they are of only two kinds: (a) "final" moves lead to an end of the game with a well specified outcome that each player values more or less; and (b) "non-final" moves lead to another single node. Such a description of a game is called the "extensive form."

1.2.1 The Extensive Form

The simplest non-trivial extensive form game is *Game 1*, illustrated in Figure 1.1.

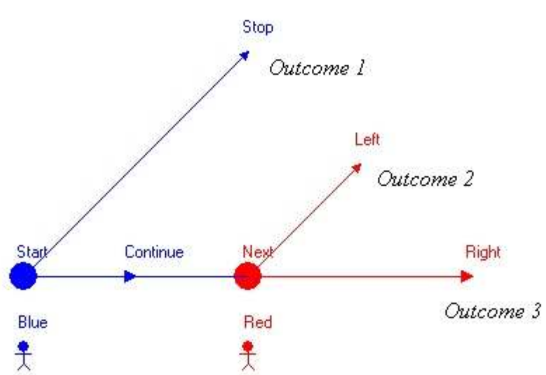

Figure 1.1: The simplest game

The player called Blue here moves first, at the very start of the game, and may choose Stop which ends the game right away in Outcome 1, with Red not even getting a chance to do anything. But Blue may choose Continue that leads to the node Next where Red has the choice between Left and Right. Each of these last two moves is also final and ends the game in one of the other two possible outcomes.

In Game Theory each outcome is valued by a numerical "payoff" or "utility" for each of the players. The higher the utility, the more desirable the corresponding outcome is for the corresponding player. In Figure 1.2, *Game 1* has utilities associated to each final outcome with their values preceded by "U=" and colored according to what player they are assigned to.

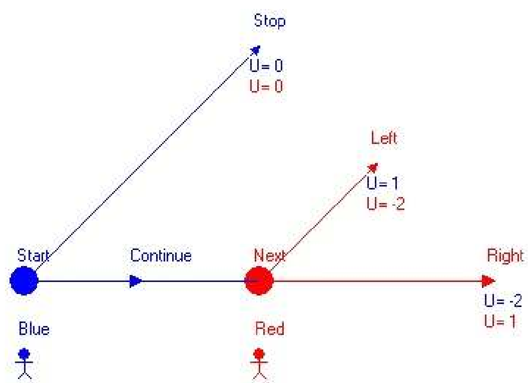

Figure 1.2: The simplest game with payoffs

All players of a game are assumed to have "perfect knowledge" about the game they are playing. This means that the entire structure of *Game 1* is understood by both sides. Moreover, all players are assumed rational so that they will always choose the maximum expected payoff given their current information and expectations. So, if the node Next is reached, Red should rationally choose Right which yields the payoff U=1 for Red, higher than U=-2 from the choice of Left. Formally, one says that Right is a "best reply" to Continue.[7]

With this reasonable expectation of Red's intentions, Blue can expect U=-2 from the choice Continue and should therefore rationally choose Stop for the higher payoff U=0. One says that Stop is a best reply to Right. But if he now plans on choosing Stop, Red never gets a chance to make a move, although she can still be planning one, just in case. But Right is still the best she can do (although secretly planning Left would have no consequences.) The choices Stop for Blue at Start and Right for Red at Next, although the latter is never implemented in this reasoning, being best reply to each other together form what is called a "Nash equilibrium".[8] Neither side can benefit from unilaterally deviating from their assigned choice. Indeed, Blue would definitely loose by choosing Continue, eventually getting a payoff of U=-2, while Red would see no improvement in planning to choose Left since she would still get U=0 from Blue's choice of Stop.

In general, players may have more than two options at their turn of play and more than one turn. The game of Figure 1.3 is called the Centipede and it can be lengthened at will to give each side as many turns as desired.

---

[7] Some authors prefer using the term "best response."

[8] It is even a subgame perfect equilibrium (SPE) that will be discussed later.

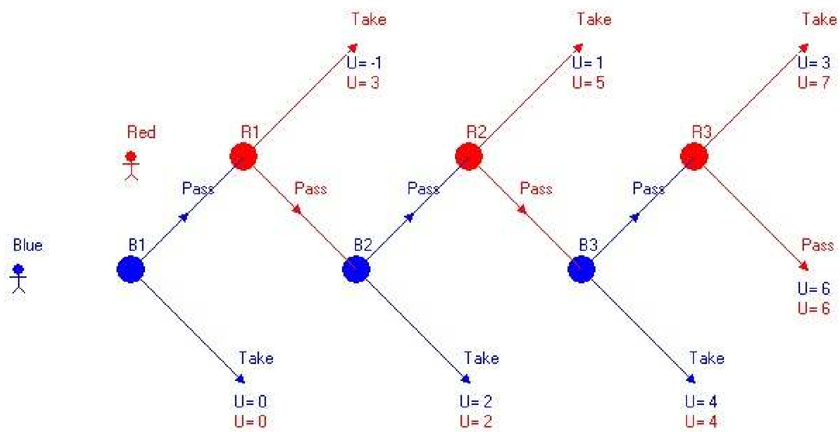

Figure 1.3: The Centipede game

Here, each side still has only two possible moves at each of his or her turn. But since, in Game Theory, a strategy means *a complete contingency plan*, Blue must consider all possible combinations of moves at each of his nodes. He may for instance plan Pass at nodes B1 and B2 and Take at B3. But he may instead plan Pass at B1 and Take at each of B2 and B3. One could reasonably argue that making any plan at B3 when he is planning Take at B2 is irrelevant. Indeed, regardless of what Red does, Blue can never reach B3 with such a plan. However, a complete contingency plan means planning even for what is not expected to happen. In this case one sees that there are 8 distinct strategies for each side in Figure 1.3. The reason for considering even unlikely circumstances can be seen in the example of Figure 1.4. This is still Game 1 but with slightly changed payoffs.

The previous line of reasoning still applies: if Red reaches Next she should choose Right, which is her best reply to Continue. And Blue should choose Continue, his best reply to Right, thus yielding a Nash equilibrium. But there is another odd and troubling plan for Red: could she threaten to choose Left if she ever reaches Next? If he believes that threat, Blue finds it best to choose Stop to the benefit of Red. And Red can, indeed, plan to respond “optimally” to Stop by Left since that does not affect her outcome. So, the strategy pair {Stop, Left} does form a Nash equilibrium, albeit not a very convincing one. In the terms that we will explore later, it is not subgame perfect. In everyday terms, one would deem Red’s threat not credible since she would hurt herself by executing it should her bluff be called by Blue with the choice of Continue.

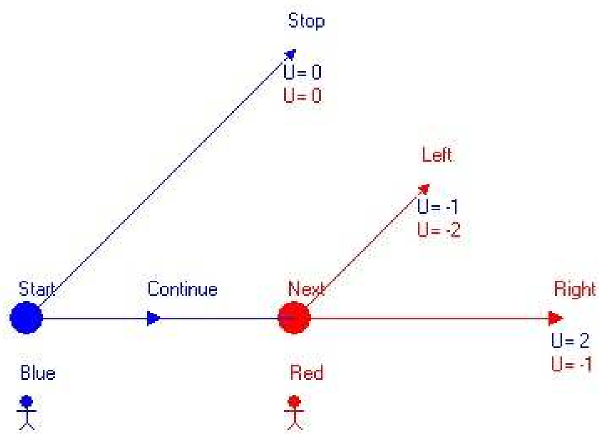

Figure 1.4: Another simple game

1.2.2 Imperfect Information

From a structural standpoint, *Game 1* and the Centipede are games of "perfect information." There is no ambiguity for Red when the node Next is reached in *Game 1*: She can infer with certainty that Blue chose the move Continue and can act with certainty on the consequences of her choice. In many circumstances, however, information is not quite so perfect although there are various instances of "imperfect information." Games of that sort always involve so-called "information sets." Instead of being at just one node at her turn of play, Red may instead be at one of several nodes but she might not be quite sure which one. In *Game 2* pictured in Figure 1.5, Red has an information set made up of the two nodes Next 1 and Next 2, and represented by a thick dotted line joining the two.[9]

At her turn of play, Red is now uncertain about what Blue did: if he chose Up, she is at node Next 1, but if he chose Down she is at Next 2. Because of her uncertainty, Red is unsure about the consequences of her own decision: if she is at Next 1, Left is better than Right. But if she is at Next 2, Left is worse than Right. This situation would arise in a game where the two sides move simultaneously, or secretly, and their respective choices are only revealed once they have both made their decisions. It is essential that Red has the same available moves (Left and Right) at each node of her information set. Indeed, if that was not the case she should be able to tell where she is by simply looking at her available moves.

But one can apply to *Game 2* the same thinking that led to the Nash equilibrium of *Game 1*: should Blue expect Red to choose Left, he finds that his best reply is to choose Up since this yields payoff U=1 rather than U=-1 if he chose Down. And if Red expects Up, she should

---

[9] It might be better to call this a "disinformation set" since it represents a lack of information. In the extensive form a "turn" will mean an information set or a single node. Asingle node forms a trivial information set called a "singleton."

choose her best reply Left for a payoff U=1 rather than Right for U=0. The pair {Up, Left} therefore provides a Nash equilibrium of this game.

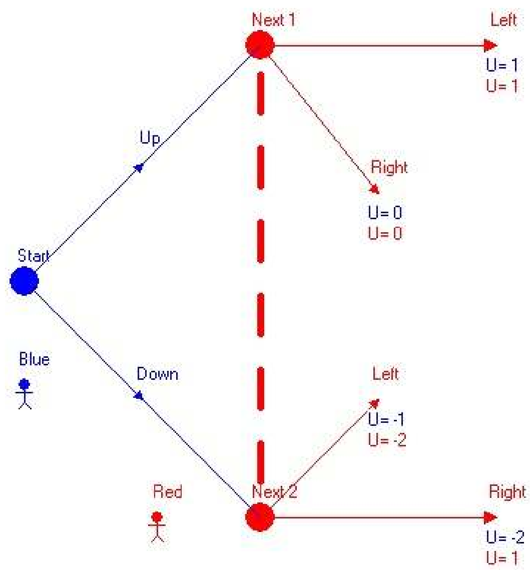

Figure 1.5: An imperfect information game

However, suppose that *Game 2* is edited with payoffs for Blue after {Up, Left} and {Down, Left} simply exchanged. The result looks like Figure 1.6:

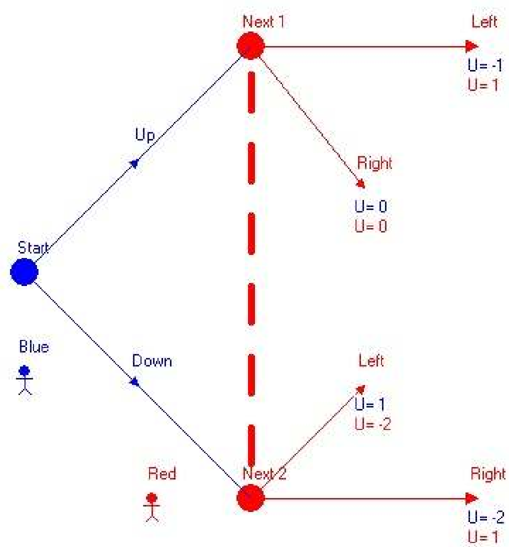

Figure 1.6: An Edited *Game 2*

Now, the choice of Left by Red calls for the best reply Down by Blue, which calls for the best reply Right by Red, which calls for the best reply Up by Blue, which leads Red to prefer Left… back to the beginning! This cyclical thinking has no end but it illustrates a fundamental idea of Game Theory: rational players can think about what the others are rationally thinking about what they are rationally thinking, indefinitely if that is necessary… Unfortunately, this best reply to best reply to best reply thinking does not lead to the identification of a Nash equilibrium in the edited *Game 2*. But this does not mean that rational players cannot reach any conclusions from this iterative thought process. Nash's insight, perhaps inspired by von Neumann's for zero-sum games, was to consider a set of plans wider than the basic Up or Down, and Left or Right: he also considered those involving probabilistic choices.[10]

There is actually a single pair of strategies (the Nash equilibrium for the edited *Game 2*) that escapes the above circular logic and provides best replies for both sides. But it specifies that Blue should pick Up with probability p=¾ and Down with probability p=¼ while Red should choose Left and Right with equal probability p=½. Indeed, we will find that each plan is a best reply to the other in the sense of expected payoffs, thus avoiding the endless cycling outlined above.[11] However, this raises interpretation issues: first, what is the meaning of making a choice with probability? Should Red just flip a fair coin in order to decide what to do? Or should the probability be interpreted as a frequency of play? This could make sense if the game is played repeatedly. Indeed, the well known "Heads and Tails" is best played by maintaining an equal frequency of either side of the coin in order to avoid being too predictable and exploitable. But not all games are played repeatedly. Besides, repeated games raise other issues mentioned before on interpreting the past and appraising the future that have not been discussed yet.

Imperfect information arises in situations other than the quasi-simultaneous play implicit in *Game 2*. In *Game 3* illustrated in Figure 1.7 uncertainty is about who moves first: Here, there is a node called Start colored in gray. That color is reserved for "Nature" turns in these notes. Any move from such a node must involve specified probabilities. Here, the probabilities p=½ of "Blue First" or "Red First" mean that Nature gives each side equal chances of being first to make a choice. However, the blue and red information sets mean that neither side knows who has that privilege when making their choice. Blue's attempt to stop the game will only succeed if he indeed wins that draw, yielding him a payoff U=1. But if Red is awarded the first move, Blue's choice of Stop will yield U=-2 when Red chooses Stop, and U=-1 when Red chooses Continue.

Game 3 is a pre-emption dilemma: if Continue is interpreted as waiting and Stop as throwing the first punch then neither side looses when both sides plan to hold off. But if either side fears the other's temptation to preempt then it is best to attempt preemption, even when one is a bit too slow in doing so. If instead of punching one considers pulling a gun the social

---

[10] Von Neumann's result cannot be applied to *Game 1* or *Game 2* since they are not zero-sum games.

[11] Blue's expected payoff of choosing Up is ½×(-1)+ ½×(0)=- ½ while his expected payoff for Down is ½×(1)+ ½×(-2)=- ½. So, either choice yields the same payoff and is therefore trivially optimal. And so are the choices of Up and Down with probability. The argument is symmetric for Red.

result can be worse. And if instead of pulling a gun one considers conducting a nuclear first strike then the future of humankind is at stakes.

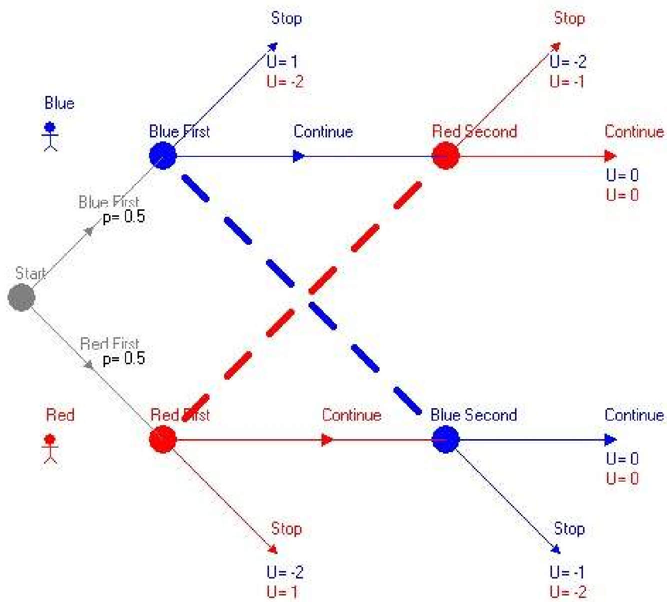

Figure 1.7: Who moves first?

Red can investigate the rationality of each choice just as in the previous games. If Blue plans to choose Stop then Red would expect U=-2 at either of the nodes Blue First and Blue Second. So, her choice Continue yields the expected payoff U=-2 given her assumption that Blue chooses Stop. If she chooses Stop under that same assumption, she instead gets U=1 at node Red First. Since that node is reached with probability p=½ she expects E=½×(1)+ ½×(-2)=- ½ from her own choice of Stop. Therefore, she does best by responding with Stop to her expectation of Stop by Blue. Clearly, since the game is entirely symmetrical, Blue will respond optimally to Red's Stop with his own Stop. So, the pair {Stop, Stop} is *the* Nash equilibrium of that game.

Well, not quite yet… It is indeed *a* Nash equilibrium of *Game 3*. But what if Red instead assumed that Blue chooses Continue? It turns out that she then has the same expected payoff U=0 from either of her two moves. By symmetry, so does Blue. So, {Continue, Continue} is also a perfectly valid Nash equilibrium of *Game 3*, although it does not seem quite as solid as the other one since unilateral deviation by either side is harmless to the deviator. Yet, *Game 3* illustrates again the fact that a Nash equilibrium is not necessarily, and is indeed rarely unique. It always exists but sometimes in excess.

1.2.3 Incomplete Information

Imperfect information can take a specialized form called "incomplete information." The idea was introduced by Harsanyi to represent situations where one side is unsure about the other players' exact priorities or capabilities. In essence, one describes the game entirely as if the players were completely informed. Then one duplicates that game and adds information sets to represent one of the player's uncertainties. The process can be replicated at will to depict two-sided (or more) incomplete information.

Game 4 illustrated in Figure 1.8 is perhaps the simplest example. The top part of the game involving nodes Honest/Blue and Honest/Red, together with the attached moves Stop, Continue, Risk and Safe is very similar to *Game 1*. It is replicated in the lower part of the game with slightly changed payoffs. Then an information set is added that represents Red's uncertainty about where she is, should her turn to play materialize. To complete the edifice, a Chance node with two possible moves and probabilities specify the upfront chances that Red will be in the upper or the lower part of the game.

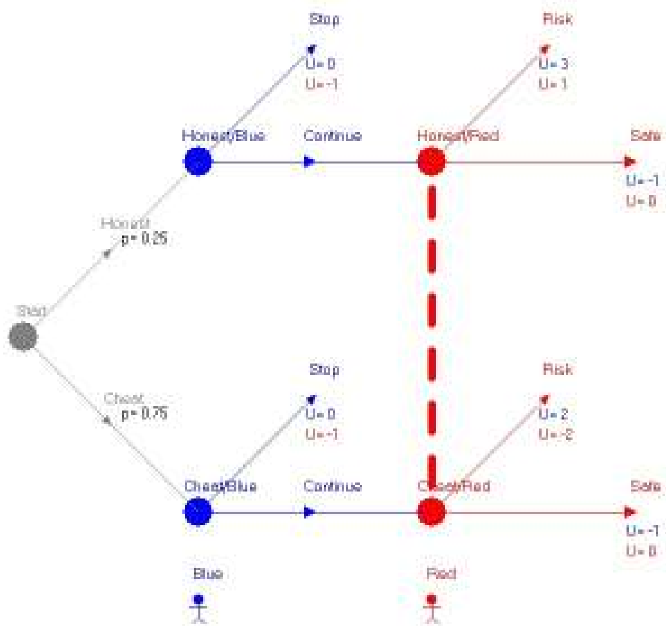

Figure 1.8: An incomplete information game

The uncertainty here is not just about Blue's priorities. It also involves the consequences for Red. They are reflected in Red's payoffs according to her moves. Of course, she has the same available moves at her two nodes; otherwise she would know where she is. In real life, the situation depicted by *Game 4* can occur when a Blue player makes an offer to a Red player who doesn't know whether it is too good to be true. Blue could be honest, but he could be a cheat. If he is indeed a cheat and she takes the risk then she loses (U=-2.) If she is risk-averse in some ways, she can play Safe, which means keeping some distance or taking some insurance that minimizes both the benefits and the potential damage. But expecting Red to play Safe discourages Blue to make the offer in the first place, and this is detrimental to Red who prefers to turn down an offer than getting no offer since she always gets the payoff U=-1 when Blue chooses Stop.

Blue knows who he is since there is no information set joining the two blue nodes.[12] So, he can make plans in the usual manner: if he expects Red to play Safe, his payoff will be U=-1, regardless of who he is, when choosing Continue. With that expectation, He should therefore choose Stop for a payoff of U=0. If, instead, he expects Red to choose Risk then he should pick Continue, no matter who he is. But how does Red assess her options? A simple approach is to weigh the chances of Blue being a cheat according to the *known* probabilities at Start.[13] When choosing Risk, she would get U=1 with probability ¼ and U=-2 with probability ¾. She therefore expects E=¼×(1)+¾×(-2)=-5/4 from the choice of Risk, worse than the U=0 she can guarantee herself by playing it safe. With that rational plan, Blue should choose Stop and this provides, indeed, a Nash equilibrium of that game.

However, the above calculation is based on a naïve appraisal by Red of her chances of being at either of her two nodes. In essence, she believes that she is at node Honest/Red with the initial probability p= ¼ that Nature selected the honest "type" of Blue. This ignores the possibility that an honest type might not behave exactly the same way as a cheat. If that is so, the probability of reaching node Honest/Red might not be quite the same as the initial chance of facing an honest Blue. We will investigate this possibility later with a discussion of Bayesian updating of beliefs.

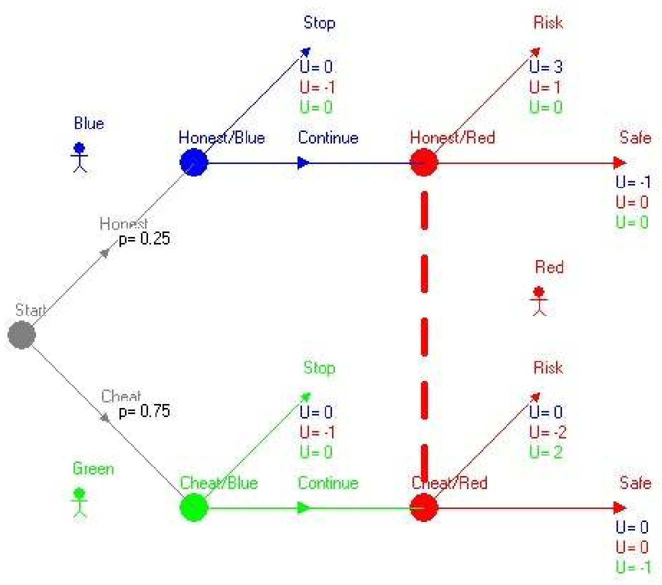

Figure 1.9: A 3-Way version of Game 4

The game of Figure 1.9 has most features in common with Game 4. The main difference is that instead of having two types of Blue there are now two distinct players Blue and Green.

---

[12] Indeed, one can replace the single player Blue by two players, say Blue and Green, each controlling one of the two nodes. See Exercise… in the Homework section.

[13] Perfect knowledge means that the Chance moves' probabilities are known to the players.

Note the indifference of Blue to all outcomes in the lower part section where he has no role to play and symmetrically for Green in the upper part. The two games have in fact the exact same solutions so that the differences are really a matter of interpretation.

1.2.4 The Normal Form

In Game Theory a strategy is a complete contingency plan for a player. This means that the strategy will specify what a player intends to do at each of his/her turn. A turn may mean either a single node (singleton) or an information set. The plan must even specify what the player will do at turns that are not expected to be reached. For instance, in the Centipede of Figure 1.3 Blue cannot just specify that he will Take at node B1 without expressing what he intends to do should nodes B2 or B3 be reached, even by accident. A strategy is called “pure” if the choices it specifies are with certainty rather than with probabilisty.

There are clearly 2^3=8 pure strategies for each player in the game of Figure 1.3. Once each player has picked a pure strategy the unfolding of the game is unambiguous. For instance, if Blue adopts the strategy {Pass, Pass, Pass} while Red adopt {Pass, Take, Pass}, the game will go through four turns and terminate at node R2 in the outcome valued (1,5). Evidently, each pairing of pure strategies yields a similar game play and eventual outcome. The “normal form,” also called the “strategic form,” of a game is a listing of all players’ pure strategies together with the anticipated payoffs for all matching of strategies. The normal form is usually presented in table form in the case of two players. The game of Figure 1.2 therefore has the normal form shown in Figure 1.10.

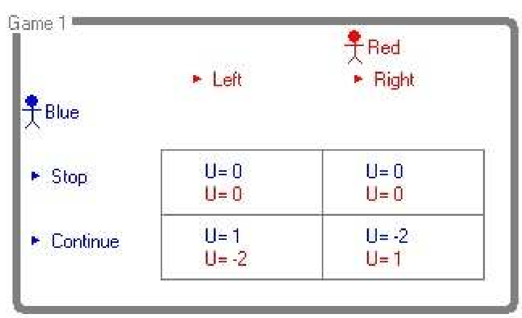

Figure 1.10: The Normal Form of *Game 1*

The reasoning on the extensive form can now be revisited with the normal form: if Red chooses Right we are in the second column of the table of Figure 1.9. Blue should respond with Stop, and that selects the upper-right cell of that table. Red, then, has no further incentive to move since Left would only yield the very same payoff U=0. So, the Nash equilibrium is, as expected, the pair {Stop, Right}. Indeed, starting from any of the four cells in the table of Figure 1.9 leads to the same stable outcome from which neither side would wish to deviate.

The only notable difference is that the above reasoning is obtained in the strategic form of the game. In the normal form, each player implicitly *commits* to a plan before the game starts instead of waiting his or her turn to make a move. This seems innocuous in the initial Game 1. But what about the edited version of Figure 1.3? The normal form now looks like Figure 1.11:

Game 1

Red

Blue

	Left	Right
Stop	U= 0 U= 0	U= 0 U= 0
Continue	U= -1 U= -2	U= 2 U= -1

Figure 1.11: The Edited *Game 1*

The usual reasoning applies: should Blue choose Continue, Red optimally responds with Right which yields the best reply Continue for Red and a Nash equilibrium. But should Red consider Left, Blue would best respond with Stop which would still yield Left as a best reply for Red, the other Nash equilibrium previously identified. The problem is that the above normal form does not even begin to suggest a possible flaw that is clearly apparent in the extensive form: Red will know what Blue did in the extensive form and may very well change her mind about implementing her threat of going Left should Blue pick Continue. In the above normal form, the timing of the moves is completely erased and Red is committed to the Left strategy before the game even starts. Indeed, Figure 1.11 is not the normal form of just the extensive form game of Figure 1.3. It is also that of the extensive form of Figure 1.12:

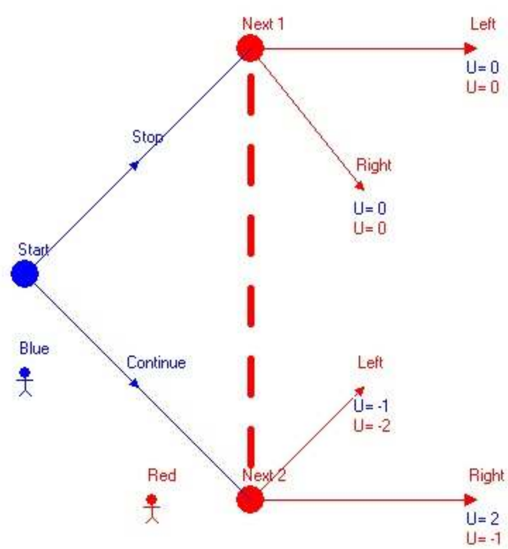

Figure 1.12: Another Extensive Form for Figure 1.11

The above extensive form is in fact a game that is very distinct from that of Figure 1.3. Here, Red has no knowledge of what Blue did when her turn of play comes. Her decision will in fact involve beliefs about what should be expected of Blue. But the important point is this: an extensive form game has a single normal form representation, up to permutations of players and strategies. The converse is not true: a given normal form can correspond to various extensive forms with different structures.

The edited *Game 2* of Figure 1.6 has the normal form shown in Figure 1.13.

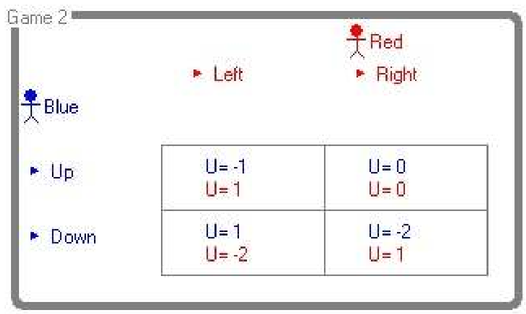

Figure 1.13: The Normal Form of *Game 2*

The reasoning about Figure 1.6 applies to the normal form in Figure 1.13: the choice of (column) Left by Red calls for the best reply (row) Down by Blue, which calls for the best reply Right by Red, which calls for the best reply Up by Blue, which leads Red to prefer Left, and so on in counterclockwise fashion. There is no Nash equilibrium made up of pure strategies in this game.

John Nash's insight, perhaps in no small part inspired by John von Neumann's earlier work, was to consider the wider set of "mixed strategies" in the search for equilibrium. A mixed strategy is simply a probability distribution over the set of pure strategies. A Nash equilibrium is in general a strategy profile, meaning a set of one strategy per player, such that each player's (pure or mixed) own strategy in the profile is a best reply to the others' in the sense of *expected* payoffs. *Game 2* admits the Nash equilibrium pictured in Figure 1.14. Red is assumed to use Left or Right with equal probability $p=\frac{1}{2}$ while Blue will use Up with probability $p=\frac{3}{4}$ and Down with probability $p=\frac{1}{4}$.

To understand how this forms a pair of best reply strategies one needs to calculate the expected payoffs to each side. The expected payoff for Blue of using Up is $E=\frac{1}{2}\times(-1)+\frac{1}{2}\times(0)=-\frac{1}{2}$. Similarly, the expected payoff of using Down is $E=\frac{1}{2}\times(1)+\frac{1}{2}\times(-2)=-\frac{1}{2}$. So, either Up or Down is trivially a best reply to Red's mixed strategy and so is Blue's mixed specified strategy in the sense of expected payoff. The reasoning holds similarly for Red's strategy with expected payoff $E=\frac{1}{4}$. The beauty of Nash's finding spelled out in the next section is that such an equilibrium always exists.

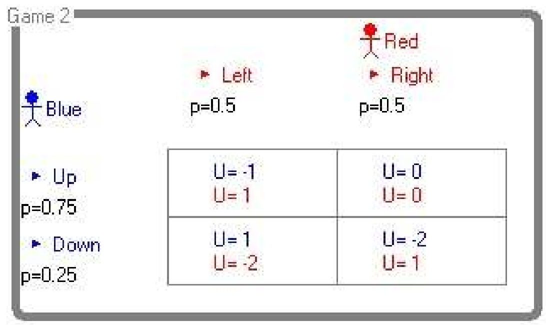

Figure 1.14: A Mixed-Strategy Nash Equilibrium

From an interpretive viewpoint, mixed strategies suggest that players use a random device to decide at the last moment (according to the given probabilities) what pure strategy to really use. This is hardly convincing. A better approach is to reinterpret the mixing in terms of the corresponding probabilities of each move in the corresponding extensive form. Then a given move might be "threatened" with the given probability, a more appealing interpretation.[14] There is an even more persuasive interpretation of probabilistic play that was introduced by Harsanyi and that we will discuss in a later chapter.[15]

The normal form of Game 3 shown in Figure 1.16 can be obtained with slightly more work. The payoffs of lower-left cell {Continue, Stop}, for instance, can be obtained as follows: if Nature chooses Blue First, the path of play ends in payoffs U=2 for Blue and U=-1 for Red. If instead Nature chooses Red First, the play ends immediately in payoffs U=-2 for Blue and U=1 for Red. In expectation, Blue receives E=-2 while Red receives E ½×(-1)+½×(1)=0. The other cells are obtained similarly. One easily finds in that table the two Nash equilibria {Stop, Stop} and {Continue, Continue} that were previously obtained.

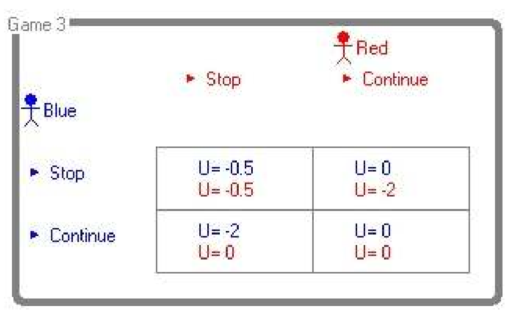

[14] This is only possible if the mixed strategy does translate into such probabilistic moves, something that requires the assumption of "perfect recall."

[15] The idea called "purification" is that players are always a little uncertain about their counterparts and the probability of facing one type of player or another can explain the probability of that player making one move or another.

Figure 1.15: The normal form of *Game 3*

The process becomes a bit more complex for *Game 4*. First, Red still has two strategies Up and Down. But what about Blue? Two types of Blue do not mean two *distinct* players. A good way to think about a game of incomplete information is like when playing poker: you are the very same person but you may or may not have an ace in the hole. Nature deals you that card at the last moment and you have to make a choice to fold (Stop) or to ante up (Continue). If you adopt strategies upfront, as is implicit in the normal form, you make up your mind about what you would do in the two cases. So, you have a plan if you find yourself at node Honest/Blue and a plan at node Cheat/Blue, each of which you implement once the missing card is dealt to you. So, there are *four* distinct strategies for Blue:

1. Stop at both nodes (denoted S|Ho & S|Ch);
2. Stop at Honest/Blue and Continue at Cheat/Blue (denoted S|Ho & C|Ch);
3. Continue at Honest/Blue and Stop at Cheat/Blue (denoted C|Ho & S|Ch);
4. Continue at both nodes (denoted (C|Ho & C|Ch).

The normal form is thus a 4 by 2 table as in Figure 1.16.

Game 4

Blue \ Red	Risk	Safe
S|Ho & S|Ch	U= 0 U= -1	U= 0 U= -1
S|Ho & C|Ch	U= 1.5 U= -1.75	U= -0.75 U= -0.25
C|Ho & S|Ch	U= 0.75 U= -0.5	U= -0.25 U= -0.75
C|Ho & C|Ch	U= 2.25 U= -1.25	U= -1 U= 0

Figure 1.16: The normal form of *Game 4*

In order to understand the payoff entries it is essential to understand the meaning of strategy: since it is a complete contingency plan made before the game is actually played, it ought to be assessed versus similar plans also made before the game is played. It is as if one looks at the Start node in Figure 1.5 when constructing the table in Figure 1.15. Viewed from that point, strategy "S|Ho & C|Ch" for Blue, paired with strategy Risk for Red (1st column, 2nd row in the payoff table) yields two possible paths: if Nature picks the Honest type the game immediately ends in Stop with the payoff pair (0,-1). But if Nature picks Cheat the game unfolds through Continue to Cheat/Red and ends with Risk in the payoff pair (2,-2). Expected payoffs are therefore: ¼×(0,-1)+¾ ×(2,-2)=(3/2,-5/8) as recorded in the corresponding cell.

The standard analysis can now be applied: starting with "S|Ho & S|Ch" Red may pick Safe from which Blue would never want to deviate since all his other payoffs in the second column are lower. And Red has no incentive to deviate from Safe in the first row. This is an already identified Nash equilibrium. Another one is displayed in Figure 1.17.

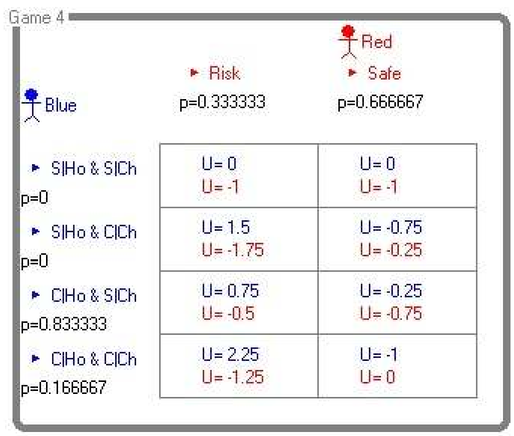

Figure 1.17: A Nash equilibrium of Game 4 in normal form

This is a mixed-strategy equilibrium for Game 4. It formally means that Blue is planning to use the two pure strategies "C|Ho & S|Ch" and "C|Ho & C|Ch" with respective probabilities p=5/6≈0.833333 and p=1/6≈0.166667. Of course, since they both specify the adoption of C|Ho, this means that Blue will play Continue with certainty at node Honest/Blue. And clearly, he will play Stop with p=5/6 and Continue with p=1/6 at node Cheat/Blue. Red's mixed strategy is interpreted similarly and corresponds to the intention of choosing Risk with probability p=1/3 and Safe with probability p=2/3 when her turn comes. There are in fact other equilibria in Game 4, but their discussion will be postponed to a later chapter. However, it worth looking at the extensive form version of the above equilibrium, that is shown in Figure 1.18.

As expected, Red chooses Risk with probability p=1/3 and Safe with probability p=2/3 when her turn comes. And indeed, Blue chooses Continue with certainty when Honest and with probability p=1/6≈0.166667 when a Cheat. So, the mixed strategies of Figure 1.17 do translate into probabilities on moves. In the extensive form, one doesn't call such probability assignments a "mixed strategy," a name that is reserved to the normal form. Instead, it is often called a "behavioral strategy."

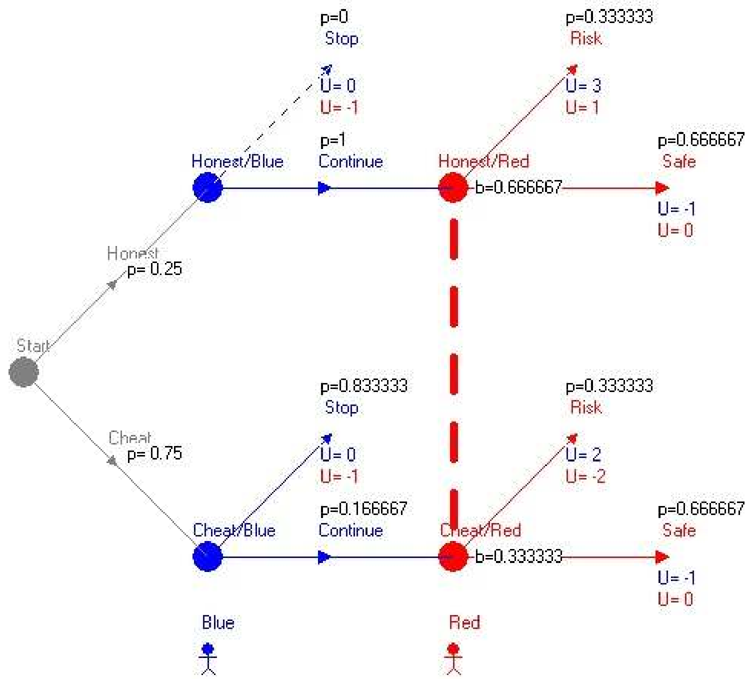

Figure 1.18: The Extensive Form Version of the Equilibrium of Figure 1.17

But something new appears in Figure 1.18 that deserves much attention: there are now numbers b=0.666667 and b=0.333333 attached to the two nodes of the red information set. These stand for "beliefs:" they mean that when her turn comes, Red has some probabilistic assessment of what node she is at in her information set. How Red can form such beliefs is quite easy to understand from the picture: the probability that play will reach node Honest/Red is simply p=¼×1=¼ given the chances that Nature picks Honest and the certain choice of Continue by the Honest Blue. And the probability that node Cheat/Red is reached is p=¾×(1/6)=1/8, given the chances that Nature picks Cheat and the probability p=1/6 that a Cheat Blue will choose Continue. Therefore, the probability that the red information set is reached adds up to p=1/4+1/8=3/8.[16] So, if Red finds herself reaching her turn she can deduce from the laws of conditional probabilities

P(@node Honest/Red)=P(Blue Honest|Red turn)×P(Red turn)

Or P(Honest|Red turn)= P(@node Honest/Red)/ P(Red turn)=(1/4)/(3/8)=2/3≈0.666667

The probability P(Honest|Red turn)=b=2/3 is Red's belief that she is at node Honest/Red at her turn of play. The belief at node Cheat/Red is obtained similarly and is the complement b=1/3. The process just outlined is called "Bayesian updating" of beliefs and will be formalized later. But it is important in order to understand the logic of the above equilibrium: only because

[16] Indeed, the probability that it is not reached is the probability that the move Stop at Cheat/Blue is made. That is precisely (3/4)×(5/6)=5/8.

of these beliefs can Red legitimately consider playing Risk and Safe with the given probabilities in the extensive form. Indeed, her choice Risk involves the expected payoff $E_B(Risk)=(2/3)\times(1)+(1/3)\times(-2)=0$ while her choice Safe yields $E_B(Safe)=(2/3)\times(0)+(1/3)\times(0)=0$. Both moves are therefore optimal given her beliefs and so is her given probabilistic play.

The equilibrium displayed in Figure 1.18 is in fact a so-called Perfect Bayesian Equilibrium (PBE) that will be studied in more details in a later chapter.

The normal form of the 3-Way version of Game 4 requires the addition of a player and can no longer be displayed in a 2-dimensional table. Instead, we may display two tables, each one corresponding to one specific choice by Green. A possible rendering is shown in Figure 1.19.

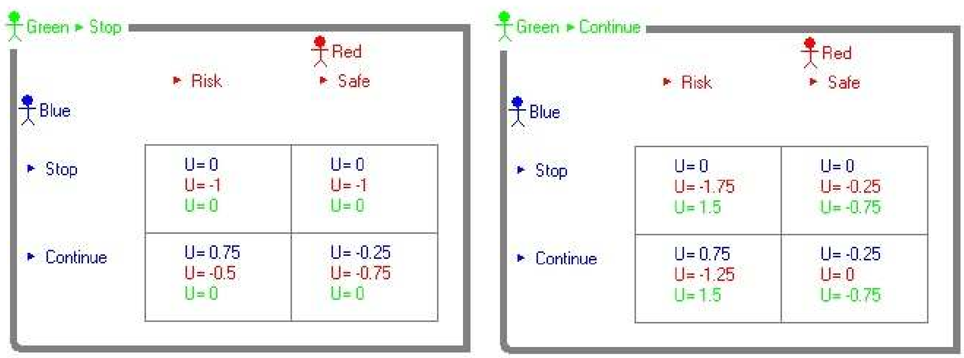

Figure 1.19: The Normal Form of the Game of Figure 1.9

A search for equilibrium can proceed in the usual way. Starting with Green Stop one applies the standard search of equilibrium between Blue and Red in the left-hand side table. This yields the stable {Stop, Safe} cell. But assuming {Stop, Safe}, Green has no interest in switching to Continue that would decrease his payoff from U=0 to U=-0.75. So, {Stop, Safe, Stop} is a Nash equilibrium of that game. There are other mixed-strategy Nash equilibria in that game that can be matched to those of the original *Game 4*.

1.2.5 The Graph Form

The games studied so far have one common feature: they can only last a few turns specified by their very structure. Of course, each such game could be player over and over again, even by the very same players. Would this affect their rational play? It depends on whether the repetition is made an explicit part of the game structure. In way of doing so is to define the game on a graph rather than on a tree so that it is not bound to end when all possible moves have been exhausted. The simplest possible example is *Game 5* illustrated in Figure 1.20.

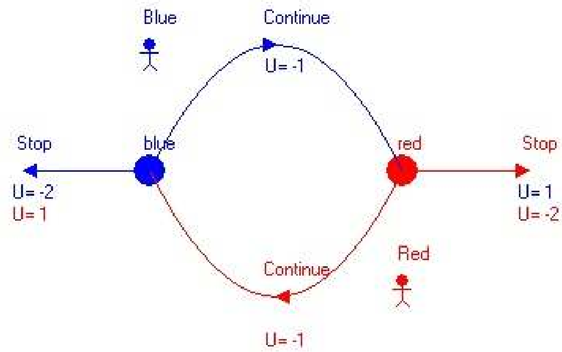

Figure 1.20: A Game on a Graph

The repeated deterministic choice of Continue by the two sides would make the game last forever. Note that the moves Continue have payoffs attached to them although they are not final.[17] The interpretation of this structure is that each side may either accept a relatively bad deal (U=−2) or decide to "pass the buck" at an immediate cost (U=−1), perhaps in the hope that the other side will accept a bad deal. One difficulty is that there is only one defined state for each side: it is either Blue's or Red's turn. Each side could of course keep track of the prior developments (the number of turns already played) but that history is not represented in any way in Figure 1.17. Representing it would require adding nodes for each side (such as "1 turn played," "2 turns played," etc.) So, there is no memory of prior developments represented in the structure of *Game 5* and there are only two pure strategies for each side here: Stop or Continue. An added difficulty is that the pair C&C={Continue, Continue} does not seem to yield a clearly defined payoff for either side. Indeed, should neither side plan to ever Stop, they would each accumulate the (U=−1) payoff indefinitely, yielding an expected payoff E=−∞. There is a simple remedy to this last issue: discounting the future. Let us now attach a discount factor, say d=0.99, to each move Continue. This really means that when contemplating a future expected payoff E_B(C&C|@red) at node red, Blue discounts it by that factor d when viewing it from node blue. And when at node red, Blue discounts the future expected payoff E_B(C&C|@blue) similarly. Now, the indefinite choice of Continue by both sides has a clearly defined payoff for Blue. When at node blue, for instance, player Blue anticipates:

$$E_B(C\&C|@blue)=-1+0.99\times E_B(C\&C|@red)=-1+0.99\times 0.99\times E_B(C\&C|@blue)$$

Solving this equation for E_B(C&C|@blue) yields:

$$E_B(C\&C|@blue)=-1/(1-0.99^2)\approx -50.25$$

[17] This feature is not specific to games on graphs although it is natural in this context. Payoffs can just as well be attached to non-final moves in the tree-based extensive form.

The reasoning is symmetrical for Red.[18] But one should note that the expected payoffs for player Blue of the strategy pair C&C is not independent of the current node (blue or red) considered. So, one cannot directly write a normal form of *Game 5* and use these expected payoffs directly. It only becomes possible if one specifies a Start node, for instance blue, and appraises all outcomes from there. In particular, the same calculation as above for player Red yields:

$$E_R(C\&C|@blue)=0.99\times E_R(C\&C|@red)=0.99\times(-1+0.99\times E_R(C\&C|@blue)$$

Solving this equation for $E_R(C\&C|@blue)$ yields:

$$E_R(C\&C|@blue)=-\ 0.99/(1\text{-}0.99^2)\approx-49.75$$

The normal form then reads:

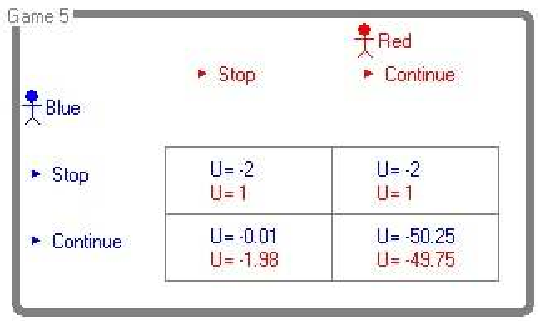

Figure 1.21: The Normal Form of Game 5 Starting at Node blue

One can now look for a Nash equilibrium as usual in the normal form (see Homework.) But one can just as well seek an equilibrium directly in the graph form. There are in fact several equilibria (see Homework 1.4.6), one of which is pictured in Figure 1.2.

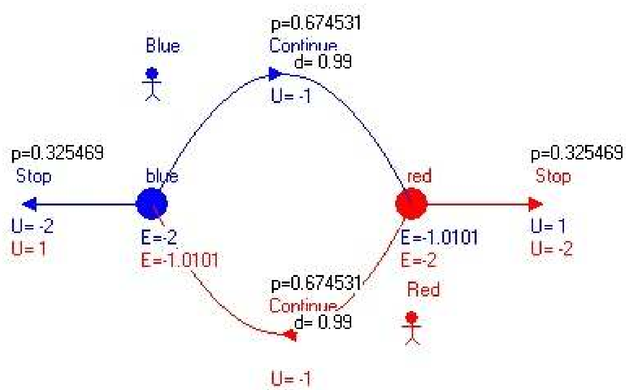

[18] The two sides need not use the same discount factor and repeating moves need not have the same discount factors either.

Figure 1.22: An Equilibrium of Game 5

At node blue, player Blue expects the immediate payoff $U=-1$ plus his discounted expected payoff of reaching node red. Figure 1.19 indicates that $E_B(@red)\approx-1.0101$ in the given profile. That expectation arises from the following calculation:

$$E_B(@red)=0.325469\times(1)+0.674531\times0.99\times E_B(@blue)$$

with the given probabilities of play by Red, and assuming that $E_B(@blue)=-2$. In turn, this last expectation results from the expected payoff for Blue of choosing Continue which yields:

$$E_B(Continue|@blue)=-1+0.99\times(-1.0101)\approx-2$$

or the same final payoff $U=-2$ for Blue of choosing Stop at node blue. Since the two moves yield the same expected payoff they are both optimal in response to Red's strategy and so is Blue's strategy of using them with the given probabilities. Of course, the argument is symmetric for Red and this indeed forms an equilibrium.[19] It is worth noting that play will continue for any arbitrarily large number of turns with correspondingly small but positive probability.

Games on graphs will be extremely valuable when modeling the repetition of so-called "constituent games," meaning games with finitely many turns that are played over and over again by the same players, with finite memory of the past and with discounting of the future. A famous example is the repeated Prisoner's Dilemma. The story goes as follows:

The district attorney has two suspects in custody but has only scant evidence of a serious crime. He places the two suspects in separates cells and talks to each of them privately, offering each the following bargain: if you confess and your accomplice does not, I will let you go with a month in jail for a minor offense. And I will use your testimony to put your accomplice away for ten years. I will offer him the same bargain, and if he confesses but you don't, he will only get a month and I will put you away for ten years. If you both confess you will each get five years. If neither of you confesses I will use trumped-up charges to get you one year each in prison.

There are many situations that share a similar structure and the name Prisoner's Dilemma is more generally associated to the normal form shown in Figure 1.23. Each of two sides has two choices: Cooperate (Coop) or Defect (Dfct). Cooperate yields a relatively desirable outcome to both sides (one one year in prison.) A unilateral Defect yields the most desirable outcome for the defecting side (one month in jail) and the worst for the cooperating side (ten years.) Simultaneous defection yields a comparatively bad outcome, not the worst, but not as good as joint cooperation. It is easy to see that Defect is a dominant strategy for each side and

[19] It is in fact a Markov Perfect Equilibrium (MPE) that we will discuss at length in a later chapter.

that the rational play of the game yields the bad outcome of the lower right-hand cell. The dilemma is that joint cooperation would yield a much better outcome, but seems unattainable.

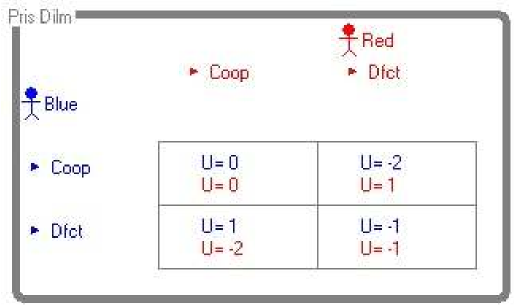

Figure 1.23: The Prisoner's Dilemma

But in many real world problems that have a Prisoner's Dilemma structure, the decision problem is encountered over and over again. Does this modify the rational outcome? It depends on whether the game is repeated a known number of times (see homework 1.4.3.) If the players don't know how long the game will last they may expect a next turn with some probability, say p=0.99. This probability can then serve as a discount factor in the graph version of the game represented in Figure 1.24.

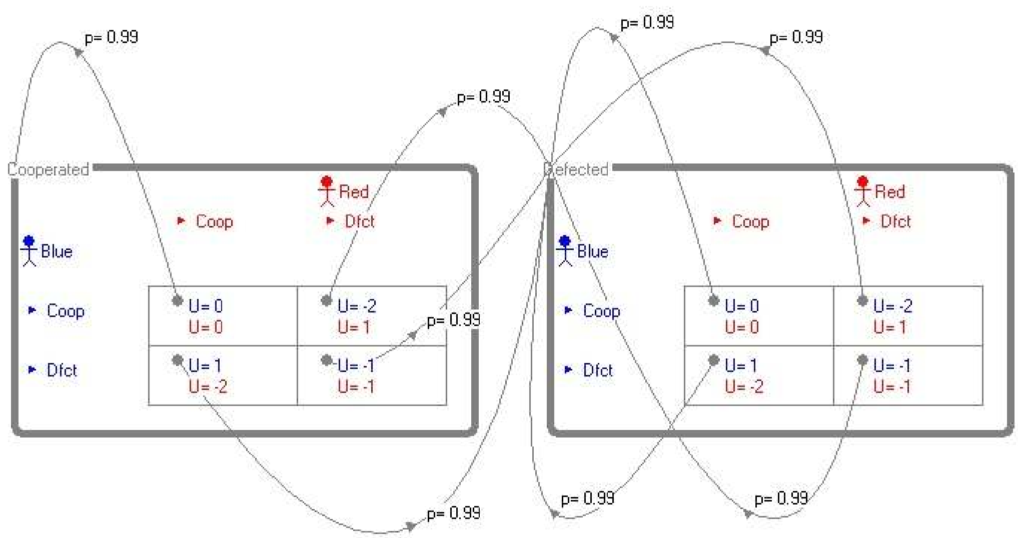

Figure 1.24: A Repeated Prisoner's Dilemma with Two Memory States

This is in fact a graph form just like Game 5, except that there are more possible and simultaneous choices available to the sides. The grey arrows indicate that, after each turn of the game, there will be a new turn with probability p=0.99. The outcome of the next play will therefore be discounted by that same factor. But that new turn may begin in either of the two tables, depending on prior developments. The left-hand side table labeled "Cooperated" can only be returned to after the simultaneous choice of Cooperate by the two sides. Any other choice or starting cell leads to the right-hand side table labeled "Defected." This means that the two sides jointly entertain only two possible memory states: either there has always been cooperation between them or there has been at least one defection by either side or both.

There are only two pure strategy equilibria in this graph form: (a) both sides always defect wherever they are; and (b) both sides always defect when starting in the Defected table and cooperate when starting in the Cooperated table. That second equilibrium is known as the "Grim Trigger." It allows cooperation to last forever but will never forgive the slightest deviation. Adding memory states can profoundly affect that picture and yield extremely interesting rational play that will discussed in Chapter 5.

1.2.6 A Continuum of Choices

Although mixed and behavioral strategies formally extend the finite number of choices to a continuous set, the nature of that extension is still severely limited: the action set for a player is a polyhedron[20] and the expected payoff is the mathematical expectation of using the probability, meaning it is multilinear. But what if the action space for a player is a more general continuum and the payoff is given by a non-linear function?

The earliest example of such a game is Cournot's duopoly. In Cournot's model, the two well owners face a demand function that expresses the price p that the local market will sustain for a global quantity Q of the good offered for sale (water in his original work) according to: $p=f(Q)$. Each of the two owners ($i=1,2$) chooses a quantity q_i to offer for sale for the day. The total supply determines a common price $p=f(q_1+q_2)$. Player i therefore receives a revenue $q_i \times p$. Costs of operation are neglected in this simple model. For a simple example, let us assume that the demand function f is given by:

$f(Q)=60-Q$ for $0 \le Q \le 60$,

and $f(Q)=0$ for $Q \ge 60$.

Given an arbitrary q_j by his counterpart, Player i wishes to maximize $U_i=q_i \times (60-q_i-q_j)$. Differentiating with respect to q_i yields the payoff maximizing:

$$q_i=30-q_j/2=\phi_i(q_j)$$

Cournot called ϕ_i Player i's "reaction function" (a reaction to Player j's choice q_j). The situation is symmetrical for Player j. When the two sides respond to each other over and over again according to their reaction function they converge to a single point $q_i=q_j=20$ where neither side can do any better, given the other side's choice. This is a standard Nash equilibrium.

1.3 A General Solution Concept

In summary, the basic *concepts* of Game Theory are:

---

[20] Formally, it is a simplex: the convex closure of the set of pure strategies.

- *Player*: an independent decision maker with influence and interests in the unfolding of a game;
- *Turn*: a node or an information set (i.e., a collection of nodes with matching available decisions at all nodes) at which a given player has a decision to make;
- *Move*: a decision available to a player at one of his/her turns. A move can lead to another player's turn or to an end of the game with a specified outcome;
- *Payoff (or Utility)*: a score on an arbitrary scale for each player that represent how each assesses one outcome versus another. The higher the score, the more desirable the outcome is for the corresponding player;
- *Extensive Form*: the specification of all players, turns, moves and payoffs of a game. Any final move must be endowed with payoffs for all players and there must be a single distinguished Start node at which the game is assumed to begin;
- *Strategy*: the specification for a player of what he/she will do at each of his/her possible turns of the game. Pure strategies stipulate deterministic decisions and mixed strategies are randomizations of pure strategies;
- *Belief*: a probability distribution over the nodes of an information set;
- *Strategy Profile*: the specification of one strategy (pure or mixed) for each player;
- *Expected Payoff*: the mathematical expectation of payoff given a strategy profile;
- *Normal (or Strategic) Form*: A list of players, pure strategies and expected payoffs for all players and strategies of the game;
- *Best Reply*: a strategy for a player that maximizes his/her expected payoff in response to all other players' strategies in a given profile;[21]
- *Nash Equilibrium*: a strategy profile such that each player's assigned strategy in the profile is a best reply to the other players' assigned strategies in the profile;

The fundamental result of game theory is:

Theorem: Any game in normal form (with finitely many players and pure strategies) admits a Nash equilibrium.

For his findings, John Nash was awarded the Nobel Prize in Economics. His proof of the above theorem involves elegant but advanced mathematics.

---

[21] The player's own strategy in the profile is not relevant to the best reply calculation.

1.4 Homework

In all the following exercises you may use the GamePlan software (see Chapter 2.)

1.4.1 Matching Pennies

Each of two players Blue and Red chooses privately to turn a penny to heads or tails. Then they reveal their choices simultaneously. If the pennies match (i.e. both show heads, or both show tails) Blue keeps them both. If they don't match (one heads and one tails) Red keeps them both. Write the normal form of this game and identify a Nash equilibrium. Justify your answer by calculating expected payoffs.

1.4.2 The Centipede

Write the normal form of the Centipede shown in Figure 1.3 and identify all of its Nash equilibria in pure strategies.

1.4.3 The Finitely Repeated Prisoner's Dilemma

Assume that the one-shot Prisoner's Dilemma is repeated finitely many times (say ten times, for example.) Could that affect the two sides' strategic thinking and yield a cooperative outcome? Hint: on the tenth and last game, how should rational players play expecting no future strategic interaction? How should they thus play on the ninth turn?

1.4.4 The Battle of the Sexes

The story, as told by Luce and Raiffa in their famous text *Games and Decisions* goes as follows (p. 91): "A man (He), and a woman (She) each have two choices for an evening's entertainment. Each can either go to a prizefight or to a ballet. Following the usual cultural stereotype, the man much prefers the fight and the woman the ballet; however, to both it is more important that they go out together..."

Assume that both sides' worst outcome is when they go their separate ways each to the other's favorite event. The next worst is when they go their separate ways, but each to their own favorite event. The best outcome, for each, is when their partner joins them in their favorite event and their next best is when they join their partner in his or her favorite event. Construct a normal form of the game, choosing carefully the payoffs to reflect the stated priorities, and find all pure Nash equilibria.

1.4.5 Selten's Horse

The following extensive form game is attributed to Selten.

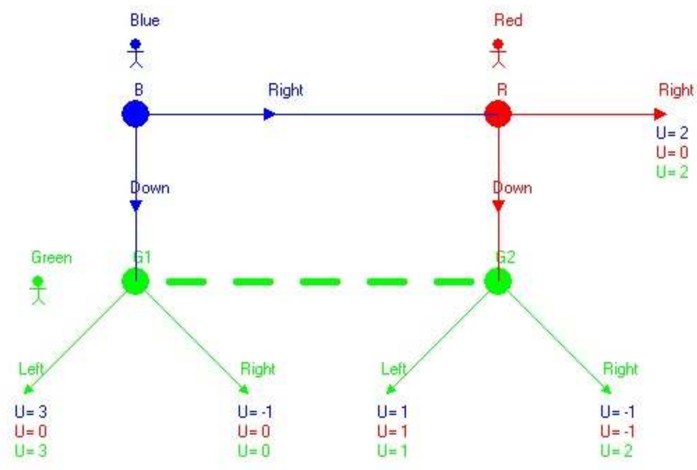

Figure 1.25: Selten's Horse

Find two pure equilibria of this game. Discuss what the corresponding beliefs should be in the green information set.

1.4.6 Critical Beliefs

GamePlan shows two equilibria for Game 4. One is shown in Figure 1.18 and the other one is shown in Figure 1.25 below.

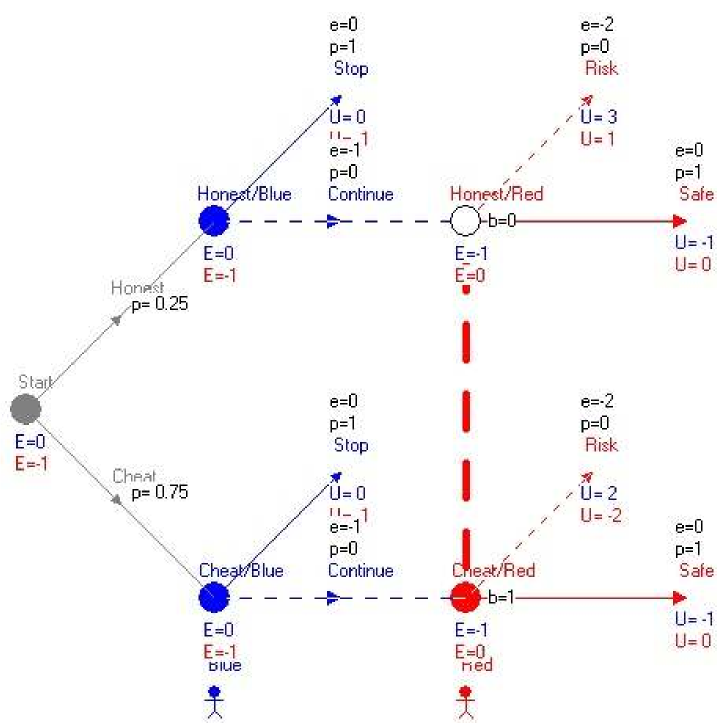

Figure 1.26: A Second Solution of Game 4

(a) Explain in your own words why these are completely rational plans for the two sides given the beliefs shown at the red nodes, and why these beliefs are justified given the plans;

(b) Show that beliefs b=0.5 at both red nodes are still compatible with the same plans and that the same plans are still optimal given the beliefs.

(c) Obtain a value b such that belief is b at node Honest/Red and belief is (1-b) at node Cheat/Red with the following property: the resulting expected payoffs of Safe and Risk are the same for Red. Such beliefs are called "critical."

1.4.7 The Simplest Graph Form

Find two pure Nash equilibria for the normal form of Game 5 pictured in Figure 1.21. Find the matching pure equilibria in the graph form.

1.4.8 Mixed and Behavioral Strategies

In the extensive game form pictured in Figure 1.24 Blue has two turns. The corresponding normal form is on the right in the picture.

Consider the behavioral strategy where Blue plays with probability p=½ each of his two available moves at each of his two turns. What mixed strategy does that correspond to in the normal form? Consider the mixed strategy where Blue plays each of his pure strategy except Up&Right with equal probability p=1/3. What behavioral strategy does that correspond to in the extensive form? Verify that this mixed strategy forms a Nash equilibrium when played against the mixed strategy with probability p= ½ on each of Stop and Continue by Red. What do you conclude?

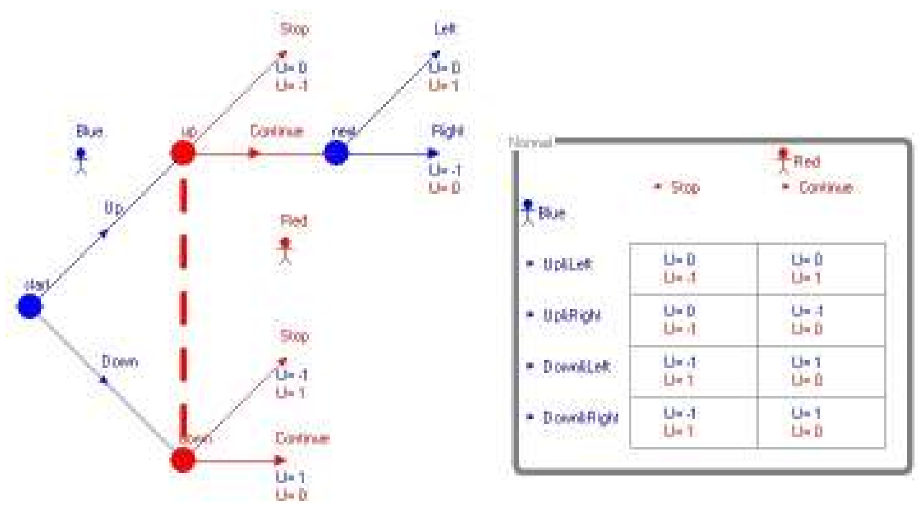

Figure 1.27: Mixed and Behavioral Strategies

It can be argued reasonably that pure strategies Down&Left and Down&Right should not be even listed since they have no reality in the extensive form.[22] Rewrite the above normal form with the resulting three meaningful pure strategies for Blue, solve and comment.

1.4.9 Three Friends around a Table

Three friends are sitting around a table in a restaurant and want to experiment with the following game: An envelope containing a $1 and a $2 bill will be passed clockwise. Each person can either take the envelope or pass it to their left. Whoever takes the envelope ends the game, keeping the $1 bill and giving the $2 bill to the friend on their right. So, each player's favorite outcome (which yields $2) is to pass to one's left and see that person take the envelope. The next best (which yields $1) is to take the envelope at one's turn. And the worst (yielding $0) is to see the person on one's right take the envelope instead of passing it. The initial player will be chosen at random by the waiter (who is assumed unbiased.)

Draw a graph form representation of this game. Assume a common discount factor d=0.999 on all non-final moves and find an equilibrium inspired by that of Figure 1.22 for Game 5. Hint: by symmetry, all three players can be assumed to have the same probability p of passing. In order for that probability to be non-trivial (neither 0 nor 1) their expected payoff of passing should be equal to their payoff of taking. *Show* that this probability p must then satisfy the equation

$$1=d\times[2(1-p)+d^2p^2]$$

And solve for p given d=0.999. Check your solution on your graph form by constructing expected payoffs at all three nodes following the example of Game 5.

1.4.10 The Cournot Oligopoly

Extend the analysis to the case of three well owners with the very same demand function as in section 1.2.6. How do prices and equilibrium payoffs for each well owner evolve as a third competitor enters the market? What do you expect would happen as the number of entries increases?

---

[22] The same point can be made about the extensive listing of strategies in the Centipede: if one player chooses Take at his or her first turn there is really no real need to consider what that player may do at later turns. In that perspective, there are really only four relevant pure strategies for each side in that game: {Take, Pass&Take, Pass&Pass&Take, Pass&Pass&Pass}.