Harvard Business School

9-187-159

Rev. April 13, 1992

Games of Strategy: An Introduction

The purpose of this note is to introduce you to some of the basic tools we will use in this course and to serve as a background to our class discussions. It should also serve as a dictionary cum reference manual for the minimal necessary terminology of the subject.

Let us begin by asking what a "game of strategy" is. Loosely speaking, any situation where the fates of two or more persons are inexorably linked may be called a game of strategy. Now clearly, this definition is too general and all encompassing. Even if all of life is indeed a game, we have a more modest and narrower goal. Let us qualify our first clumsy attempt at a definition and try again. Any situation where the choices of two or more rational decision makers together lead to gains and losses for them is called a game. In addition, a game may simultaneously involve elements of both conflict and cooperation among the decision makers. The use of the word "game" is unfortunate. Although most parlor games like bridge, chess etc. qualify by our definition, the subject is interesting only because it deals with more serious matters also. It is hard to find a business situation that does not qualify as a "game". Biologists have used the subject to gain insight into evolutionary processes. Political scientists use the subject to study the most serious game of all the arms race. For better or worse, the term game is commonly used to describe such situations and we are stuck with it.[1]

Let us start by looking at an extremely simple, stylized problem. Suppose that as the manager of a firm you are considering the possibility of entering a completely new market in which there is just one other firm operating at present. While the new market appears profitable and you are pretty confident of your ability to make money in it, the biggest uncertainty you face is how the entrenched incumbent firm will react to your entry. It could be accommodating, allowing you to carve your own niche in the market. On the other hand, it could follow an aggressive response, meeting you with price cuts and discounts. Suppose that in the first case your profits would be $5 million and in the second case you would suffer losses totaling $7 million. The "decision analysis" approach to this problem would ask you to assess some probabilities for the uncertainty that you face the response of your rival and then ask whether entry into this new market is worthwhile. Suppose that you assess the probability of aggressive behavior on the part of your rival as one-half. You might consider drawing a decision tree that looks somewhat like the following:

---

[1]Goethe once complained: "Mathematicians are like Frenchmen: whatever you say to them they translate into their own language, and forthwith it is something entirely different."

Professor Vijay Krishna prepared this case as the basis for class discussion rather than to illustrate either effective or ineffective handling of an administrative situation.

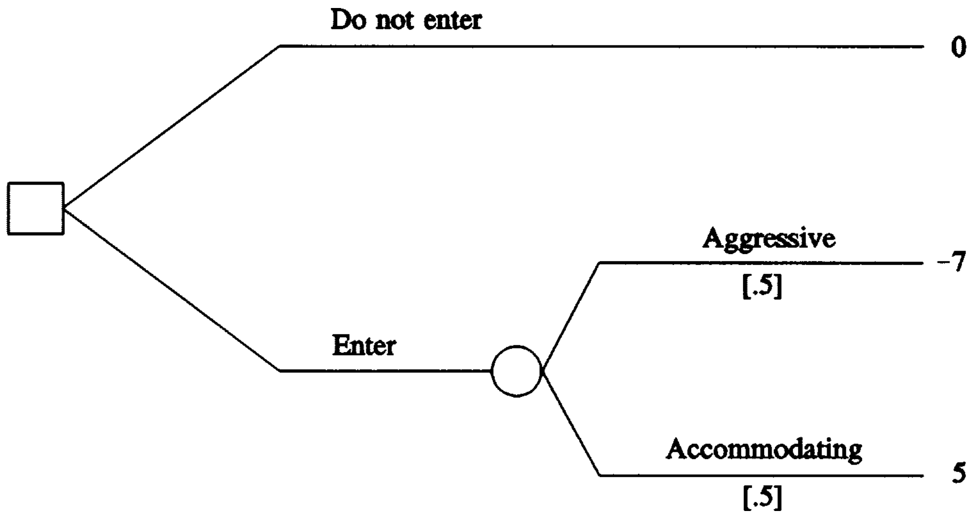

From here on, the analysis is extremely simple, you would look at the expected monetary value (EMV) of each decision. If you enter, the EMV is $[0.5×5 + 0.5×(7)] = $1 million and thus you expect losses. It does not make sense to enter the new market. Is this correct? It is, if you believe the tree we have drawn and the data it contains. The key aspect of the problem is the correct assessment of the probabilities. If the rival firm were less likely to follow an aggressive strategy in response to your entry, it could be worthwhile to enter this new market. So what are the correct probabilities? The feature that distinguishes game theory from decision analysis is that if the uncertainty you face is a result of another decision maker's actions in this case, the rival firm you should expect that the other decision maker will act in his or her own interests also. While this seems like an obvious idea, it does have far reaching consequences. Let us reexamine the simplified problem we began with again, this time adopting the game theoretic viewpoint.

From your rival's point of view, the best outcome would be if you would stay out of the market altogether. Suppose that if it is the only firm operating, the rival's profits are $10 million. If you do enter and the rival firm is accommodating, its profits are $8 million. These are higher than your profits in this event because the rival firm has been in the market for some time and probably has a loyal customer base. Finally, if you enter and the rival firm meets you with an aggressive response of price cuts and discounts (which you will have to match) its profits are a mere $2 million. With this information, let us draw a new tree, this time not a decision tree but rather what is called a *game tree*. The major difference here is that the "uncertain" behavior of the rival, which was depicted as an event node (a circle) has been replaced by a decision node (a square). To distinguish between decision makers we have identified the nodes with numbers. The entering firm is firm 1 and the rival incumbent firm is firm 2. Furthermore, as end-consequences, we write the profits of both firms. The first number denotes the profits of firm 1 and the second those of firm 2.

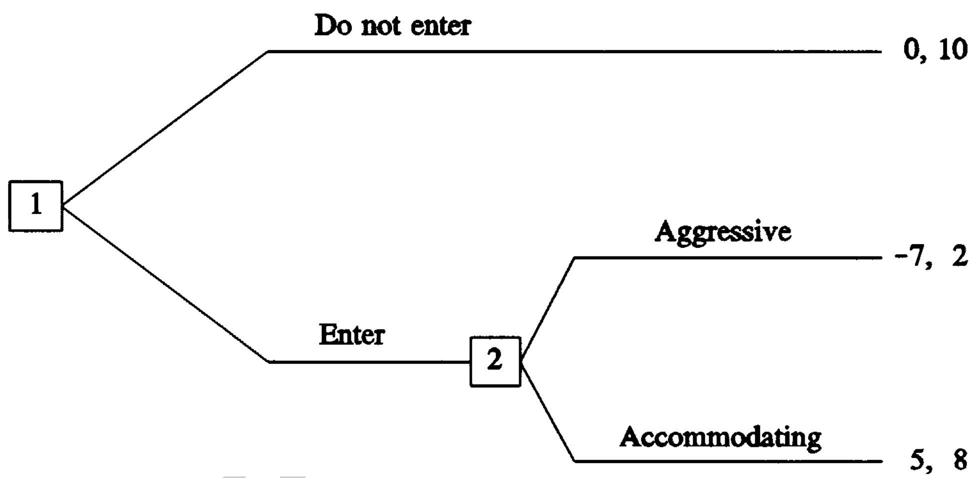

Having drawn the game tree we proceed to "fold it back", just as we would fold back a decision tree. We reason as follows. If you were to enter, it would clearly not be in the rival firm's interests to follow an aggressive strategy since this would net it only $2 million. Accommodation, on the other hand, would net it $8 million. Thus, if you were to enter, the rival would make room for you in the market. As a result, entering this new market would net you $5 million and is certainly worthwhile.

Notice that the rival's behavior is derived by considering what its interests are and thus, assuming that it too is rational and interested in maximizing profits, we are able to circumvent the problem of having to assess the probability of various actions (fight versus accommodate) the rival could take in response to your entry.

The example above was deliberately chosen to be extremely simple, but it aptly illustrates what is probably the key feature of game theory: the assumption that all actors in a strategic situation behave as rationally as you do. In other words, while thinking rationally, there is no reason to ascribe irrational behavior to your rivals.[2] To the extent that this is a good assumption, our analysis is useful. Later in this note we will return to this issue.

Game Trees

In the previous section, we saw how a simple strategic problem could be represented in the form of a game tree. It is not too surprising that *any* strategic situation can be represented in this fashion. Let us look at another, more complicated, but familiar example: chess.

In chess the first player, white, has a total of twenty possible opening moves each of the eight pawns can be moved in any of two ways (one or two squares forward) and each of the two knights can be moved in two ways (to rook-3 or bishop-3). Following a choice by white, black also has a total of twenty possible moves. A portion of the tree is given below and, as you can see, there are a lot of branches. Even after only one move each there are already four hundred (twenty times twenty) possible positions.

[2]A sage once said: "He is a fool that thinks not that another thinks".

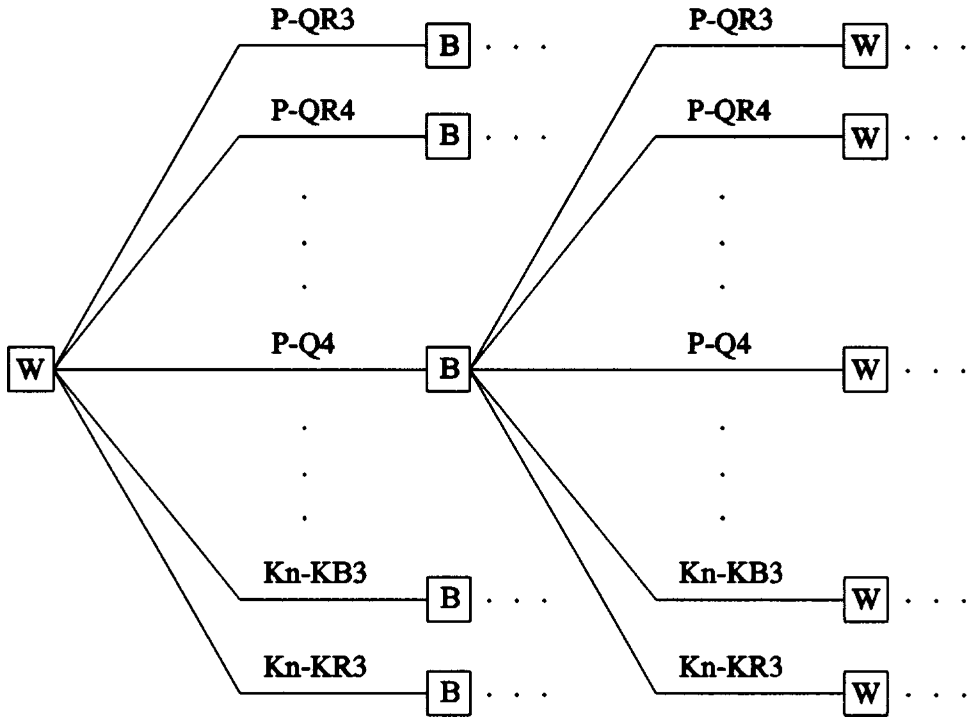

In principle, given a large enough piece of paper, enough ink, time and patience, we could draw a complete tree for chess.[3] Of course, a supercomputer would come in handy. Suppose that as a result of superhuman effort we could, in fact, draw a complete tree for chess. We could then proceed to fold it back just as we did the very simple tree in the previous section. We would, of course, arrive at one of three possible answers the result of folding back the tree would mean either that white won the game or that black won the game or that the game was a draw. Now this does not seem too surprising given that these are the only three ways in which a game of chess can end. But, in fact, folding back the tree would yield a stronger conclusion we would be able either to prescribe to white a way to win *always* (even if he or she were matched against the best grandmaster, a Kasparov or a Karpov) or to prescribe a way for black to win *always*, or to prescribe a way for both to ensure a draw every time. The point (a subtle one) is that there is actually no real uncertainty about the outcome of a game of chess provided one could draw a complete tree. The only uncertainty (and that is what makes chess interesting) is that no one has been able to draw out the complete game tree and fold it back.

It may help to think about a simpler game: tic-tac-toe. Now most of you know how to play the game so as to ensure a draw every time and that is why it is thought to be uninteresting. The complete tree for tic-tac-toe is not that complex and the method of play that always ensures a draw can be derived by folding back the tree. The point is that chess is different *only* in that it allows for a greater variety of moves and hence is more complicated. If a large enough computer becomes available to draw and fold back the tree, chess will also take on the status of tic-tac-toe.

It is pretty clear how game trees are drawn. The only thing to remember is that unlike decision trees, in game trees there are at least two decision makers and one must specify at each *node* whose move it is. This is done by putting a letter (W or B in the chess example) or number to identify the relevant player.

---

[3] An important rule in chess is that if the same position appears on the board thrice during a game, the game is declared a draw. This means that no chess game can go on forever.

Since there are numerous decision makers, and their evaluations of various end-point consequences may be different, one has to specify for each player, the value of reaching a particular end-point. We have already seen how this is done in the first example of a firm deciding whether or not to enter a new market.

In some cases, there is genuine uncertainty about certain factors. Suppose there are two players (1 and 2) and before playing a single game of chess, they toss a coin to decide who will have the white pieces. A game tree for that will look as follows:

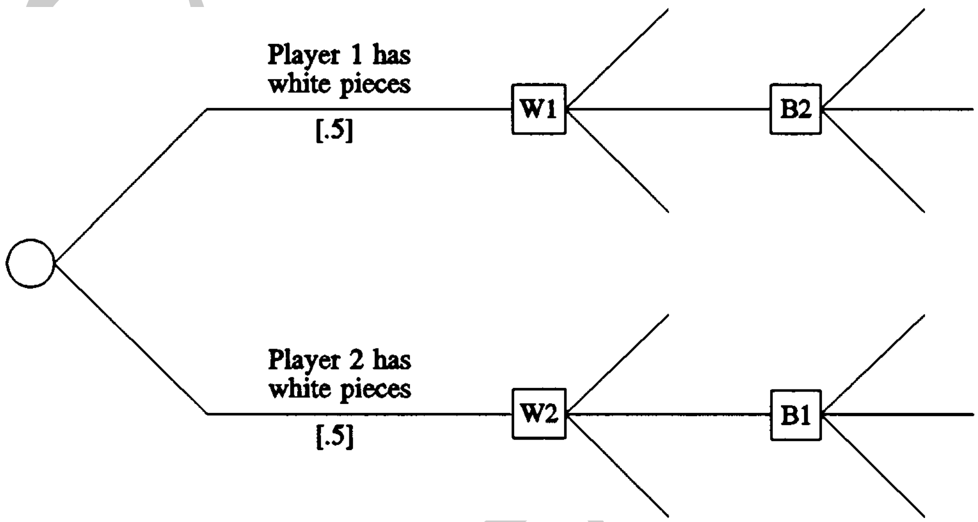

Rather than give you a list of rules about how game trees are drawn, we will learn how to do this as we proceed. The thing to remember is that they are nothing more than multi-person decision trees and it is important to keep track of which nodes involve decisions by other players and which nodes are truly chance moves (involving the toss of a coin or the vagaries of the weather, etc.).

In the sort of trees we have drawn so far, each player knows all the moves that have preceded his or her move. Again, the chess example is a good one. At any stage in the game each player observes the moves made by his or her opponent explicitly ("he played P-Q5 in response to my Kn-B4 which followed his Q-R3 ..."). In many games the play of opponents remains unknown to you. You may not know how much a rival firm is spending on advertising when you have to make a similar decision. In the course, we will see how one can represent such games by means of game trees also.

What is a Strategy?

While you are familiar with the word "strategy" in an informal sense, its widespread use in many different contexts makes it hard to define accurately. In common usage the connotation is that of a long term plan well thought out in advance. In our language a strategy will also be a long term plan that takes into account *all* contingencies one might anticipate as arising. In principle, once we have drawn a game tree to depict a particular situation, we can list all possible "strategies" for each player. Let us see exactly how this is done.

Suppose you have drawn a game tree. A strategy for a particular player, say player 1, is a complete plan of what choices that player will make at *each* node in the tree where he or she is supposed to make a choice. The idea that you have a plan ready for every node (i.e., every contingency) is crucial.

In many instances, it is impractical to write down a complete list of strategies, again because the game tree may be too large to analyze. In chess a strategy is a complete plan which can tell you how to play in any position in which you might find yourself. For instance, it might take the following form: "I will begin by playing P-Q4; if black plays P-K3, I will counter with P-K3 also; if black plays P-Q4, I will play Kn-QB3; otherwise, I will play B-K5...". As you can see, strategies in chess can get rather complicated very quickly as they involve thinking far ahead and planning responses to each of millions of possible positions in the future. Even the best of grandmasters have plans for at most six or seven moves ahead.

The idea of a strategy, while simple enough, forms the key to systematic thinking about games. The question of "how to play the game" is answered by finding the "best" strategy to adopt. In what follows, we will attempt to lay down some general principles which will aid in the selection of optimal strategies.

From Game Trees to Matrix (Strategic) Form

Suppose we are looking at a two-person game and have drawn a game tree to depict the strategic situation. Given any pair of strategies, one for each player, we can, by working through the tree, determine the end consequence that results. In our chess example, if we knew the strategy white had adopted, that is, his planned move in every contingency, and similarly knew black's strategy, we could figure out who the winner would be (or whether the game would be a draw). Of course the number of strategies a chess player can choose is enormous approximately $10^{10,000}$!

Let us put our first example, the simple entry game, in strategic form. In this game the first player has only two strategies (whether or not to enter) and the second player also has exactly two strategies (whether to be aggressive or accommodating). We can represent the game in the form of a matrix, also called a *strategic form.*

		Firm 2	
		Aggressive	**Accommodating**
Firm 1	**Do Not Enter**	**0, 10**	**0, 10**
	Enter	**-7, 2**	**5, 8**

Notice that if firm 1 decides not to enter, the outcome is the same (0, 10) no matter what strategy firm 2 adopts. This seems strange on the face of it, but recall that firm 2 has to make a decision about its pricing/marketing policy *only* if firm 1 indeed enters the market. To understand this point fully let us now look at a slightly more complicated version of the entry game.

Suppose that, as before, firm 1 is contemplating entry. Firm 2, on the other hand, has to decide its pricing policy it could keep its prices high or cut them by offering discounts. But now, more realistically, the pricing decision has to made irrespective of whether the first firm enters the market or not. A tree for this game would be different and would look as follows:

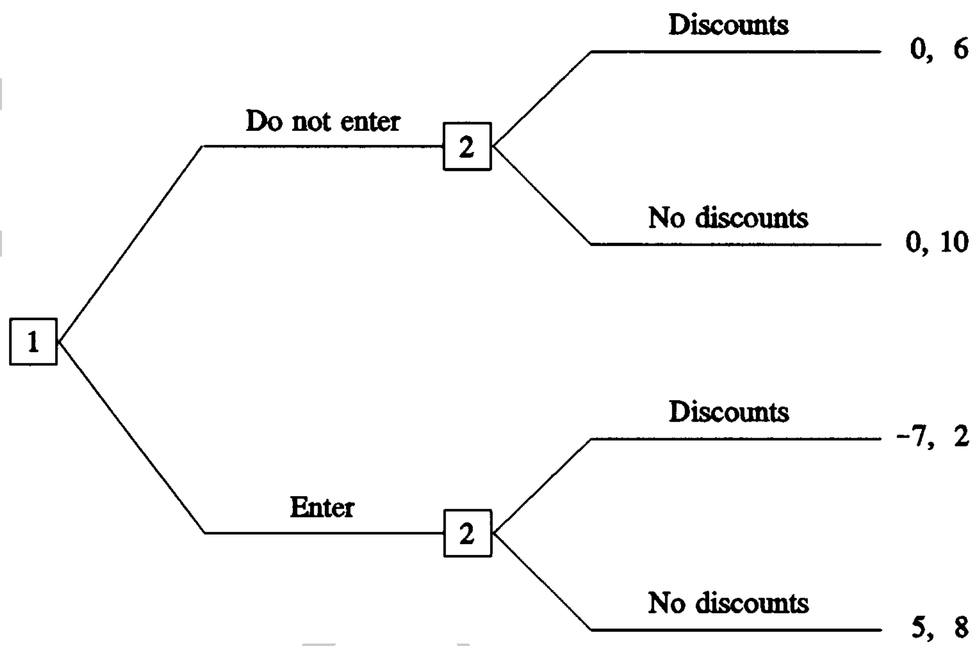

We have added a new consequence. If firm 1 does not enter and the incumbent firm offers discounts, its profits fall to $6 million. Firm 1, of course, earns nothing in this market. But notice that even though the incumbent has only two choices whether or not to offer discounts it has *four* strategies! This is because there are actually two contingencies created by firm 1's decision. We can list these strategies as follows:

(DD) Offer discounts whether firm 1 enters or not;

(NN) No discounts whether firm 1 enters or not;

(ND) No discounts if firm 1 does not enter, discounts if firm 1 does enter;

(DN) Discounts if firm 1 does not enter, no discounts if firm 1 does enter.

Now some of these strategies (DD and DN) seem silly just on the face of it. Firm 2 has no incentive to offer price discounts if firm 1 does not enter the market. Nevertheless, the strategic form looks like:

	DD	NN	ND	DN
Do Not Enter	0, 6	0, 10	0, 10	0, 6
Enter	-7, 2	5, 8	-7, 2	5, 8

Notice the manner in which the matrix is constructed. Fix a pair of strategies, say "Enter" and "ND". The "ND" strategy says that firm 2 will offer discounts only if firm 1 enters and since when this strategy is matched against firm 1's "Enter" strategy, it results in discounts being offered, the resulting outcome is that 1 loses $7 million and 2 makes only $2 million in profits. Contrast this matrix with the first one we had written down for the simpler game where firm 2 had a pricing decision to make *only* if 1 entered the market.

The difference between the matrix representation of the two games (and they are indeed different) should highlight the role of *contingent* planning. In the first game, there is no such role as firm 2's pricing in the event of no entry is fixed. In the second game above, firm 2 can decide what sort of pricing policy to follow *depending* on whether firm 1 enters the market or not.

No doubt all this seems rather pedantic at present, but you will see that care in specifying the exact game being played will pay dividends as we proceed.

We have shown that there are two different ways in which games can be represented so as to be amenable to analysis by means of game trees or by means of a strategic form (matrix). We now turn our attention to the question of finding optimal ways to play, that is, finding strategies that yield the highest expected value.

Dominated Strategies

We begin our quest for optimal strategies by asking the opposite question: when can we discard a particular strategy as clearly being sub-optimal? Suppose that there is another strategy that under no circumstances yields a lower payoff and sometimes does better. Then without fear of loss, we might as well discard the first strategy as being useless and adopt the second one. Under such circumstances, we will say that the first strategy is *dominated* by the second one. Let us look once again at the second, more complicated, entry game:

	DD	NN	ND	DN
Do Not Enter	**0, 6**	**0, 10**	**0, 10**	**0, 6**
Enter	**-7, 2**	**5, 8**	**-7, 2**	**5, 8**

From firm 2's point of view, DD, ND and DN are all *dominated* by the NN strategy. DD is clearly dominated since the payoffs from NN are strictly better whether or not firm 1 enters the market (6 vs. 10 if there is no entry and 2 vs. 8 if there is). ND is dominated because NN yields the same as ND if no entry takes place (both yield 10), whereas if entry does take place NN does better (8 vs. 2). DN is dominated because NN does better if no entry takes place (10 vs. 6) and is equally good if entry occurs (both yield 8). Thus, in this game there is no point in firm 1 playing anything but NN.

In the example above, there was a single strategy that dominated all the others. Our next example is more complex:

	X	Y
A	7, 3	0, 8
B	2, 5	3, 0
C	5, 10	0, 2
D	2, 3	5, 4

You should verify that in the example above, strategy B is dominated for player 1 by strategy D and that C is dominated by A. Thus, player 1 can safely discard both B and C. However, neither A nor D dominates the other. A does better than D if player 2 plays X (7 vs. 2) but does worse than D if player 2 plays Y (0 vs. 5). Does player 1 have some way of deciding whether A or D should be played? It all hinges on whether player 2 will play X or Y and so far there is no way to tell what he will choose. It is easy to check that from player 2's point of view, neither X nor Y dominates the other. We will see below how 1's dilemma can be resolved.

We have seen how one might check whether a particular strategy is dominated by another and argued that players can safely discard dominated strategies from their choices.

Successive Dominance

In our previous example we saw that B and C were dominated strategies from player 1's point of view and thus he would never choose either of these. Player 2, on the other hand, had no dominated strategies: X did better than Y if 1 played B or C and Y did better than X if 1 played A or D.

Is there some way for player 2 to choose between X and Y? The answer is yes. He might reason as follows: "Looking at the matrix it is clear to me that strategies B and C are dominated from 1's point of view and thus I am sure that she would never actually play either. But if she is only going to play either A or D, then Y is clearly better than X for me since it yields 8 vs. 3 if 1 plays A and 4 vs. 3 if 1 plays D. Thus it makes no sense for me to choose X."

The foregoing argument is an example of *successive dominance*. Even though 2 had no dominated strategies to begin with, once he took into account the fact that 1 would never play strategies which were dominated, X became a dominated strategy. By realizing that 1 would never play the dominated strategies B and C, 2 might as well have assumed that the following game was being played, where B and C have been discarded:

	X	Y
A	7, 3	0, 8
D	2, 3	5, 4

In this "reduced" game, X is dominated by Y and hence 2 can safely discard X. Notice that for player 1 neither A nor D dominates the other but again she could say to herself: "I know that player 2 will reason it out and realize that I will never play either B or C. Thus, he will only look at the reduced game. In the reduced game, X is dominated by Y and hence I can be sure that he will actually play Y. Knowing this, the best thing for me to do is to play D, since it yields 5 when 2 plays Y whereas A yields 0."

We have argued that in the foregoing game, player 1 should play D and player 2 should play Y. As you can see, the argument is rather complicated and relies on "reflexive thinking", that is, asking the question "what would I do if I were in the other player's shoes?" Furthermore, it relies essentially on the assumption that your opponent is as rational as you are. But we have argued earlier that there is really no reason to believe that other players are less rational than you are. Moreover, you have to believe that other players believe that you are as rational as they are, that they believe that you believe that they are rational, and so on. The important thing is the mutual recognition of the rationality of all players.[4]

Successive Dominance and Folding Back Trees

So far we have proceeded as follows. Given a game, we drew the game tree that represented the various moves and countermoves of the players. From the game tree, we wrote the "strategic form" or matrix associated with the game. We then argued that players could discard any dominated strategies and extended this principle by arguing that the discarding of dominated strategies could be done successively. This was because some strategies which were not dominated to begin with (X for player 2 in our example), became dominated once the other player discarded his or her dominated strategies (in our example once player 1 had discarded B and C). In the games that we have been considering so far,[5] this procedure usually leads to a single outcome and hence a single optimal strategy for each player.

But recall that we had an alternative method of "solving" the game by folding back the tree. So which method is right? Folding back the tree or the method of successive dominance? It would be unfortunate if the two methods gave different answers. Luckily, that is not the case and in fact there is no need to choose between the two methods. For all intents and purposes, they are the same and will yield the *same* answer.

[4]The poet Burns wrote: "O wad some pow'r the giftie gie us
To see oursels as ithers see us!"

[5]We have been considering games where every time players are asked to make a choice, they are completely aware of all previous choices of other players (as in chess). Such games are called games of *perfect information*.

Let us look into this a little more closely by looking at an example. We will do this by drawing the tree, folding it back and comparing the result to that obtained by successively eliminating dominated strategies in the corresponding matrix.

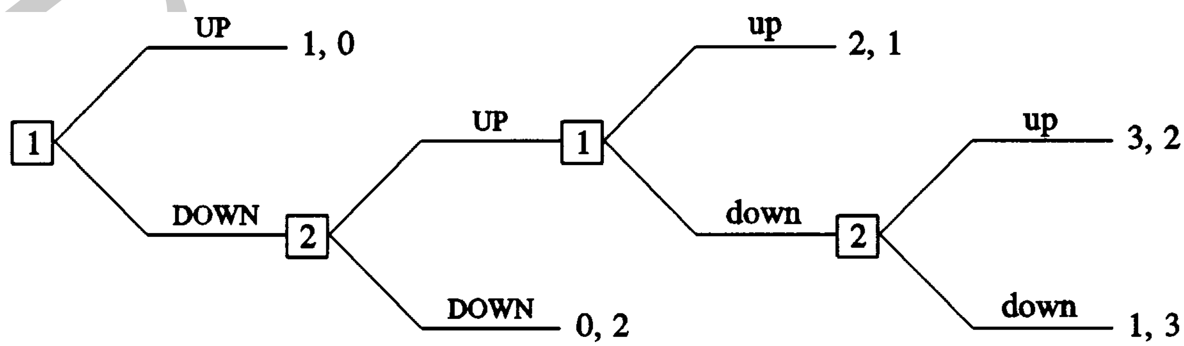

Notice that by moving "UP" at the first node player 1 can immediately end the game though doing this is quite unprofitable for both players. Similarly, player 2 can end the game by moving "DOWN" the first time he is asked to move. It is in both players' interest that the game continue as long as possible as then both will benefit. But let us see what folding back the tree yields.

If the game were to progress to the very last node where player 2 makes his second choice, he would clearly choose "down" (a payoff of 3 vs. 2). Anticipating this, player 1 would, at the middle move, choose "up" because this yields her a payoff of 2 whereas if she played "down" she is sure to get only 1. Anticipating 1's move, 2 would then choose "DOWN" in the previous move (2 vs. 1). Finally, anticipating this, player 1 would move "UP" at the very first move, since this yields her a payoff of 1 as opposed to 0. Both players want to "take the money and run" even though this is mutually harmful. The game matrix looks as follows:

	D	**Ud**	**Uu**
U	**1, 0**	**1, 0**	**1, 0**
Du	**0, 2**	**2, 1**	**2, 1**
Dd	**0, 2**	**1, 3**	**3, 2**

Notice carefully the way the game matrix is constructed. Each player has three strategies. Player 1 can, for instance, end the game immediately by playing "U" (for "UP") or she can play "Du", which stands for her choosing "DOWN" at her first move and planning to choose "up" if called upon to move a second time. Her third strategy is to play "Dd" whereupon she would choose "DOWN" at her first move and then play "down" if she does get to move a second time.

It is easy to check that player 1 has no dominated strategies to begin with. For player 2, however, "Uu" is dominated by "Ud". Once player 2 has discarded "Uu", player 1 can discard "Dd" since this now becomes dominated by "Du". In the next step, player 2 can discard "Ud" as this is now dominated by "D". Finally, player 1 now can discard "Du". Thus by successively eliminating dominated strategies, we find that player 1 will play "U" and end the game immediately.

As you can see, the method of folding back the tree and the method of successively eliminating dominated strategies both yield the same outcome. In fact you should convince yourself that throwing away a dominated strategy is exactly like taking a step backwards in the tree.

Our two methods are, in fact, equivalent and which one you use is a matter of convenience rather than of substance.

Summary

We have seen how one can analyze a simple class of games those of perfect information where each player is completely informed of all previous moves. We adopted two different techniques for solving such games and finding optimal strategies for the players. The first was to draw the associated game tree and fold it back very much in the manner in which decision trees are drawn and folded back. The second technique was to write down the strategic or matrix form and solve that by successively throwing out dominated strategies. We argued that the two procedures were in fact equivalent.

In the course, we will also talk about games of "imperfect information", that is, situations where decisions have to be made under partial ignorance of what moves other players have taken or are taking. Naturally, these situations are more complex than those we have encountered in this note. Nevertheless, many of the general techniques and principles will carry over.

Meanwhile, as Sherlock Holmes said: "Wake up, Watson! The game's afoot".