# Game Trees

Adam Brandenburger*

Version 01/04/07

## 1 Introduction

We're now going to return to non-cooperative games. So far, we have described non-cooperative games as payoff matrices. In this note, we look at another description, the game tree. Of course, we shall also have to take a look at the relationship between these two kinds of description.

## 2 Game Trees

Informally, a *game tree* (or, game in *extensive form*) is a formal description of how a non-cooperative game can be played. It specifies the moves available to each player at each point in the game, the players' knowledge when they move, and the players' evaluations of the different possible outcomes of the game (i.e. the players' payoffs).

A game tree can be thought of as a multi-person decision tree, in which the decisions of all the players are represented simultaneously. Figure 1 depicts a very simple game tree. (Here, each endpoint of the tree is associated with a pair of numbers, where the left-hand number is the payoff to player 1 and the right-hand number is the payoff to player 2.)

This game can be interpreted as one in which a player, labeled player 1, must decide whether or not to trust another player, labeled player 2. If player 1 does not trust player 2 (i.e. chooses $L$), both players receive 0. If player 1 decides to trust player 2 (i.e. chooses $R$), player 2 must react. If player 2 violates player 1's trust (i.e. chooses $l$), then both players receive a payoff of $-1$. (Think of this number as reflecting the cost of the ensuing conflict between the two players.) If player 2 honors player 1's trust (i.e. chooses $r$), then both players receive a payoff of 1.

---

*With the assistance of Amanda Friedenberg and Konrad Grabiszewski. tree-01-04-07

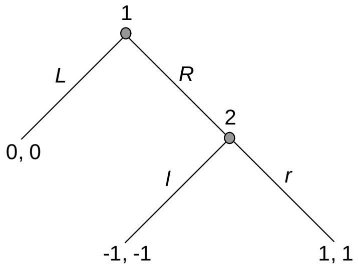

**Figure 1**

The formal definition of a game tree can be found in the classic paper "Extensive Games and the Problem of Information," by H. Kuhn, in H. Kuhn and A. Tucker (eds.), *Contributions to the Theory of Games*, Vol. II (Princeton Univ. Press 1953), 193-216. Here, we are going to proceed fairly informally. But we will now introduce a few formal terms. Thus, the black circles in the above tree are called *decision nodes*, the endpoints are referred to as *terminal nodes*, and the branches leaving the decision nodes are called *moves*.

There is one more ingredient we shall need, which is the concept of an *information set*. Consider the game tree in Figure 2 below.

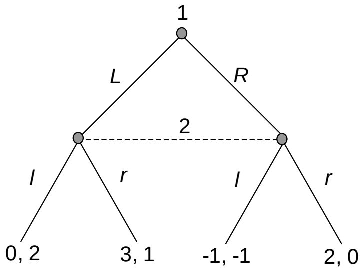

**Figure 2**

First, player 1 decides to move $L$ or $R$. Then player 2, without knowing which move player 1 made, decides to moves $l$ or $r$. More precisely, the situation is that when she has to move, player 2 does not know whether she is at her left-hand node (the one following $L$ by player 1) or her right-hand node (the one following $R$ by player 1). To represent this situation, the two decision nodes of player 2 are joined by a dotted line. A collection of nodes of this type is called an *information set* of the player in question.

Note some things about information sets. First, each of the nodes in an information set must have the same number of branches coming out of it. Otherwise, a player would be able to infer something more about which node he is at, simply by counting the number of available moves. This would contradict the idea that the player does *not* know which of the nodes in the information set he is at. Second, an obvious special case is when each information set is a singleton, i.e. consists of precisely one node. (This was the case for the tree in Figure 1.) A game is said to have *perfect information* if every information set is a singleton. Otherwise, the game is said to have *imperfect information.*

Third, suppose that instead of the situation above in which player 1 moved first, the 'physical' situation is that, first, player 2 decides to move $l$ or $r$. Then player 1, without knowing which move player 2 made, decides to move $L$ or $R$. Or, suppose the situation is that the two players make their moves simultaneously. You can see that the tree in Figure 2 is an appropriate description of *each* of these three situations. This says is that what a tree really describes is information, not time. What matters is what each player knows when he has to make a move, not whether one move is made before or after another.

## 3 Matrices vs. Trees

Since we now have two ways of describing interactions–matrices and trees–the question naturally arises as to which description to use in analyzing an underlying interaction.

Prima facie, a game matrix describes only a situation where each player makes a single choice, in ignorance of the choices made by the other players, and the game is then over. The tree thus appears to be a more general description, allowing players to move more than once and also to observe what other players do.

We are certainly going to study game trees in their own right, precisely because they explicitly model what players learn about one another's moves, and this will allow us to consider some very interesting issues that then arise.

But at a theoretical level, the extra generality of trees over matrices is not so clear. The reason is that the matrix can, in fact, be thought of as modelling *any* interaction, even ones in which the players move more than once. The key idea here is that strategies of the matrix can be thought of, not as single moves, but rather as complete plans of action for the tree.[1] This insight was one of the key contributions of Borel and von Neumann in founding game theory.

## 4 Strategies

A *strategy* of a player associates with each information set of that player, one of the moves that is available at that information set. Intuitively then, a strategy

---

[1] As the very term "strategy" suggests, which is why we used it even before introducing trees.

is a complete plan of action for the player, that specifies how the player would move in every situation where the player may have to move.

In the game tree in Figure 1, player 1 has two strategies, viz. $L$ and $R$, while player 2 also has two strategies, viz. $l$ and $r$. In the tree in Figure 2, each player again has two strategies: $L$ and $R$ for player 1, and $l$ and $r$ for player 2. Now consider the tree in Figure 3 below.

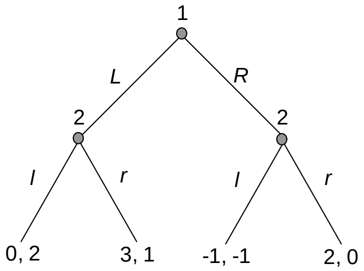

**Figure 3**

Note that this tree looks similar to that in Figure 2, but is, in fact, quite different. Unlike the tree in Figure 2, this tree has perfect information. Player 2 gets to observe player 1's move before making her own move. What are the strategies this time? Well, player 1 has the two strategies $L$ and $R$, just as before. But now player 2 has four strategies, not two. One such strategy is to move $l$ if player 1 has moved $L$, and to move $r$ if player 1 has moved $R$. Write this strategy as $(l, r)$. In the same notation, the other three strategies are $(l, l)$, $(r, l)$, and $(r, r)$. Of course, in the actual play of the game, player 1 must move either $L$ or $R$, and not both. So, only one of the two situations that player 2's strategy considers can actually arise. Nevertheless, by definition, player 2's strategy has to specify what she would do in *both* situations. (Otherwise, it would not be a 'complete plan of action.')

## 5 The Induced Matrix

With the concept of strategy in hand, we can now define the matrix naturally associated with a given game tree. Thus, starting with the tree, the first step is to compile, for each player, the set of strategies (as defined above) that are available to that player. Next, note that if for each player we specify one of that player's strategies, such a list of strategies (one per player) determines a unique path through the game tree to a unique endpoint, and hence a payoff to each player. This gives us the payoff matrix.

Going through this procedure for the tree in Figure 2 gives the matrix in Figure 4:

|  |  | 2 |  |
|---|---|---|---|
|  |  | *l* | *r* |
| 1 | *L* | 2<br>0 | 1<br>3 |
|  | *R* | -1<br>-1 | 0<br>2 |

**Figure 4**

By contrast, the procedure on the tree in Figure 3 yields the matrix in Figure 5:

|  |  |  |  | 2 |  |
|---|---|---|---|---|---|
|  |  | *l, l* | *l, r* | *r, l* | *r, r* |
| 1 | *L* | 2<br>0 | 2<br>0 | 1<br>3 | 1<br>3 |
|  | *R* | -1<br>-1 | 0<br>2 | -1<br>-1 | 0<br>2 |

**Figure 5**

Note the difference between the two matrices, which reflects the difference between the two trees they come from.

## 6 Equivalence

Now that we have seen how to obtain a matrix from a game tree, we can go back to our earlier question of which description to use. The way to make this question precise is to look more carefully at the map from trees to matrices. Specifically, does this map set up a one-to-one correspondence between trees and matrices? If it does, then the two descriptions would be completely equivalent. If not, then we will have to ponder some more.

In fact, the map does *not* set up a one-to-one correspondence. To see this, consider the two (distinct) game trees in Figures 6 and 7 below.

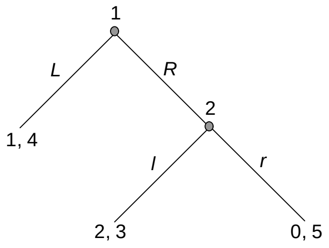

**Figure 6**

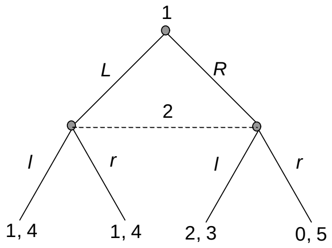

**Figure 7**

Following our procedure, both these trees are mapped to the matrix in Figure 8:

| | | 2 | |
|---|---|---|---|
| | | *l* | *r* |
| 1 | *L* | 4<br>1 | 4<br>1 |
| | *R* | 3<br>2 | 5<br>0 |

**Figure 8**

So, it seems that we would lose information about what game is really being played, if we were to conduct our analysis on the matrix and not on the tree. (We would lose the knowledge of whether it is the tree in Figure 6 or the tree in Figure 7 that is being played.) This is one reason why a large number of recent researchers in game theory choose to conduct their analyses on the tree. And, as we said above, we, too, are going to do some analysis on trees.

But this is not the end of the matter. Let us note, for specialists, where matters go next. The counter is that the two game trees in Figures 6 and 7 are, in fact, "strategically equivalent." (The idea of strategic equivalence is made precise in the papers "Equivalence of Information Patterns and Essentially Determinate Games," by N. Dalkey, in H. Kuhn and A. Tucker (eds.), *Contributions to the Theory of Games*, II, Annals of Mathematical Studies, 28, Princeton Univ. Press, 1953, 217-244; and "Equivalence of Games in Extensive Form," by F. Thompson, *Research Memorandum RM*-759, The RAND Corporation, Santa Monica, 1952.) We won't go into the details here, but we will just observe that if one does accept the strategic equivalence of the two trees,[2] then one can go further and argue that one ought to conduct analysis on the matrix and *not* on a game tree! The argument is that this way, one ensures that the analysis is not sensitive to strategically irrelevant details. (The analysis should give the same answer for the two trees in Figures 6 and 7.) For a presentation of this point of view, see the paper "On the Strategic Stability of Equilibria," by E. Kohlberg and J-F. Mertens, *Econometrica*, 54, 1986, 1003-1038.

We shall now proceed to do some analysis on game trees. This does not pre-judge the matters mentioned here. Even if one wanted in the end to define things on the matrix, there is important insight into how to do that to be gained by looking at trees.

# 7 Exercises

The following two exercises involve deciding how to go from some given 'physical' situations to appropriate trees and/or matrices.

**Example 1** *An entrepreneur is struggling with a basic marketing decision: to advertise his new product, or save the money. Cash flow is tight, and the threat of bankruptcy always seems near. The entrepreneur decides he faces the following payoffs:*

| | |
|---|---|
| *Don't Advertise & Product Sells* | 10 |
| *Advertise & Product Sells* | 9 |
| *Don't Advertise & Product Doesn't Sell* | 0 |
| *Advertise & Product Doesn't Sell* | $-1$ |

---

[2] Note for specialists: Though it turns out that to do so, one should at the same time suppose that the players avoid *weakly dominated* strategies. Weak dominance is a strengthening (!) of the concept of strong dominance from the note "Dominance and Iterated Dominance." We will come to it in "Forward Induction."

As an empirical matter, it would be very interesting to know how different 'framings' of a game, such as those in Figures 6 and 7, might influence how the players actually play.

*The right decision seems clear. Not advertising yields* 10 *or* 0*, versus* 9 *or* $-1$ *from advertising. So the entrepreneur decides to save his money and not advertise. Is something wrong here?*[3]

**Example 2** *The following scenario, termed Newcombe's Paradox, has been extensively studied in the philosophy literature.*[4] *A person is faced with two boxes: box A contains $1,000 and box B contains either zero or one million dollars. The person can choose either box B or both boxes. The prizes are placed by a genie who has profound insight into the psyche of the person and thus knows whether the person will choose both boxes or just one. If the person is to choose both boxes then the genie will put zero dollars in box B. If the person is to choose only box B, then the genie will put one million dollars into box B.*

*Should the person choose only box B, or both boxes? What is the relationship between this problem and the preceding one?*

---

[3] The problem is taken from D. Kreps, *Notes on the Theory of Choice*, Westview Press, Boulder, 1988; and the wording of this problem and the next from "Two 'Paradoxes' of Decision Theory," by Adam Brandenburger and Gus Stuart, teaching material, 2/2/99.

[4] The following rendition is from E. Dekel and F. Gul, "Rationality and Knowledge in Game Theory," in D. Kreps and K. Wallis (eds.), *Advances in Economics and Econometrics*, Cambridge Univ. Press, Cambridge, 1997.