Part II

Dynamic Games of Complete Information

10
Preliminaries

As we have seen, the normal form representation is a very general way of putting a formal structure on strategic situations, thus allowing us to analyze the game and reach some conclusions about what will result from the particular situation at hand. However, one obvious drawback of the normal form is its difficulty in capturing time. That is, there is a sense in which players' strategy sets correspond to what they can do, and how the combination of their actions affect each others payoffs, but how is the order of moves captured? More important, if there is a well defined order of moves, will this have an effect on what we would label as a reasonable prediction of our model?

Example: Sequencing the Cournot Game: Stackelberg Competition

Consider our familiar Cournot game with demand $P = 100 - q$, $q = q_1 + q_2$, and $c_i(q_i) = 0$ for $i \in \{1, 2\}$. We have already solved for the best response choices of each firm, by solving their maximum profit function *when each takes the other's quantity as fixed*, that is, taking q_j as fixed, each i solves:

$$\max_{q_i} \, (100 - q_i - q_j)q_i \, ,$$

and the best response is then,

$$q_i = \frac{100 - q_j}{2}\,.$$

Now lets change a small detail of the game, and assume that first, firm 1 will choose q_1, and before firm 2 makes its choice of q_2 it will observe the choice made by firm 1. From the fact that firm 2 maximizes its profit when q_1 is already known, it should be clear that firm 2 will follow its best response function, since it not only has a *belief,* but this belief must be correct due to the observation of q_1.

Now assuming common knowledge of rationality, what should firm 1 do? It would be rather naive to maximize its profits taking some fixed belief about q_2, since firm 1 knows *exactly* how a rational firm 2 would respond to its choice of q_1: it will choose,

$$q_2 = \frac{100 - q_1}{2}\,.$$

This, in turn, means that a rational firm 1 would replace the "fixed" q_2 in its profit function with the best response of firm 2. That is, firm 1 now chooses q_1 to solve:

$$\max_{q_1} \left[100 - q_1 - \left(\frac{100 - q_1}{2}\right)\right] q_1 \tag{10.1}$$

There is a fundamental difference between this maximization and the one done for the Cournot example, and it depends on *what the firms know when they make their choices*. Namely, in the Cournot game, neither firm *knew* what the other is choosing, so they set a belief and maximize their profit function. Here, in contrast, firm 2 *knows* what firm 1 had produced when it makes its choice, and as a result firm 1 *knows*, by common knowledge of rationality, that firm 2 will choose q_2 rationally, that is, by using its best response function. As a result, firm 1 should not use some fixed q_2 in its profit function, but the choice of q_2 that will result from any choice of q_1, which is precisely firm 2's best response function.

The solution is then given by the first order condition of (10.1), which is

$$100 - 2q_1 - 50 - q_1 = 0,$$

yielding $q_1 = 50$. Then, using firm 2's best response above we have $q_2 = 25$. The resulting profits are then $\pi_1 = (100-75)\cdot 50 = 1,250$, and $\pi_2 = (100-75)\cdot 25 = 625$.

Recall that in the original (simultaneous move) Cournot example the quantities and profits were $q_1 = q_2 = 33\frac{1}{3}$ and $\pi_1 = \pi_2 = 1,111\frac{1}{9}$. We see then that by moving first, firm 1 has a "First-Mover Advantage." (Is being such a first mover generally good? Well, would you like to be the first mover in a Matching Pennies game? Only if you want to lose...)

This example of sequential moves was introduced and analyzed by Heinrich von-Stackelberg (1934), and is known as the "Stackelberg" model. It illustrates an important point: the order of moves might, and often will matter.

In what follows in this chapter, we will set out the structure that allows us to capture such sequential strategic situations, and apply what we have already seen to these new representations. We will then go a step further and introduce a solution concept that captures an important idea of *sequential rationality*.

11

Extensive Form Games

11.1 The Extensive-Form Game

The concept and formalization of the *extensive form game* is meant to formally capture situations with sequential moves of players, and to allow for the knowledge of some players, when it is their turn to move, to depend on the previously made choices of other players. As with the normal form, two elements must to be part of any extensive form game's representation:

(EF1) Set of players, N

(EF2) Players' payoffs as a function of actions, $\{u_i(\cdot)\}_{i \in N}$

To overcome the limitations of the normal form and capture the idea of sequential play, we need to expand the rather simple notion of pure strategy sets to a more complex organization of actions. We will do this by introducing two parts to actions: First, as before, what players can do, and second, *when* they can do it. Taking the Stackelberg game as an example, we need to specify that players can choose any quantity they like, as in the normal form, but we also need to specify that player 1 moves first, and only then player 2 moves. Thus, in general we need two components to capture sequences:

(EF3) Order of moves

(EF4) Choices players have when they can move

Now that we have the players, their order of moves, and their possible actions in place, we need to introduce a notion of knowledge for players when it is their turn to move. Recall that the simultaneity of the normal form was illustrative of players who know nothing about their opponents' moves when they make their choices.

More precisely, it is not the *chronological order of play* that matters, but *what players know when they make their choices*. Using the Cournot versus Stackelberg comparison, it may be that firm 1 indeed makes its choice several hours, or days, before firm 2 does. However, if firm 2 has to make its choice *without observing* the actual choice of firm 1, then it is the Cournot model that correctly captures the situation. In contrast, if firm 1's choice is revealed to firm 2 before it makes its own choice, then the Stackelberg game is the correct representation.

In general, it may be that some information is revealed as the game proceeds, while other information is not. To represent this we need to be precise about the unfolding of information and knowledge, and for this we have a fifth component to the description of an extensive form game:

(EF5) The knowledge that players have when they can move

To add a final component, that vastly enriches our ability to describe strategic decision problems, we introduce the possibility that before certain players move some random event can happen. For example, imagine that some firm A is embarking on a research and development (R&D) project that may succeed or not. Another firm B can choose to adapt its competitive strategy given the outcome of firm A's R&D project, but may benefit from waiting to see how it develops. Since the outcome of the R&D is not fixed but rather depends on some luck, we should think of the outcomes from the R&D project as a random variable, say success or failure, each with some fixed probability.

The physical description of this scenario is as follows: first, firm A chooses its project. Second, there is uncertainty about the outcome of the project, which is represented by success with some probability p and failure with probability $1 - p$. Finally, after the resolution of uncertainty, form B will make its own choice.

We will generally call the stages where some uncertainty is present *moves of nature*. It is useful to think of nature as a player that has a fixed random strategy, and is not a strategic player. In the example above, when it is nature's choice to move it will select "success" with probability p, and this probability is fixed and *exogenous* to the game.[1] Thus, our sixth element represents nature as follows:

(EF6) A probability distribution over *exogenous* events

Finally, to be able to analyze these situations with the tools and concepts we have already been introduced to, we need to add the final and familiar requirement:

(EF7) The structure of the extensive form game represented by (EF1)-(EF6) above is common knowledge

This set up, (EF1)-(EF7), captures all of what we would expect to need in order to represent the sequential situations we want to represent. The question is, what formal notation will be used to put this all together? For this, we introduce the idea of a *game tree*.

11.2 Game Tree Representation of Extensive-Form Games

The game tree is, as its name may imply, a figurative way to represent the unfolding nature of the extensive form game. Consider, for example, the following two player game. Player 1 first chooses whether to enter (E) the game or not (N), the latter choice giving both players a payoff of zero. If player 1 plays E, player 2 can choose to cooperate, giving both players a payoff of 1, or he can defect (D) and get 2, while leaving player 1 with a payoff of (-1). A simple way to represent this may be with the graph depicted in figure 3.1

In this figure, we draw player 1 at the top with her two choices, and following the choice of N we get the payoffs determined as (0,0). Following E, however, we have player 2 moving with his two choices, and the payoffs that will be determined given each of his choices.

[1]By exogenous we mean that it is not in the choice set of players, and is fixed, just like the strategy sets and the payoff functions.

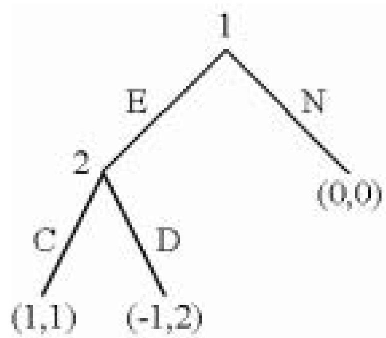

FIGURE 11.1.

This very simple structure is the most elementary form of a game tree. It includes all of the elements we described above (except for a random move by nature), but it lacks the formal structure that would clearly delineate the "rules" that are used to describe such a game tree. For example, what kind of structures are allowed? How do we capture knowledge? How do we incorporate Nature? To do this formally, there is some amount of detail and notation that needs to be introduced.[2] This is the main objective of this section, and later, once the ideas are clear, we will focus our efforts on variants of the simple game tree depicted in figure 3.1.

Definition 18 *A* game tree *is a set of nodes $x \in X$ with a precedence relation $\varphi : x > x'$ that means "x precedes x'". Every node in a game tree has only one predecessor. The precedence relation is* transitive *($x > x'$, $x' > x'' \Rightarrow x > x''$),* asymmetric *($x > x' \Rightarrow$ not $x' > x$) and* incomplete *(not every pair of nodes x, y can be ordered). There is special node called the* root *of the tree denoted by x_0 that precedes any other $x \in X$. Nodes that do not precede other nodes are called* terminal nodes, *$Z \subset X$, each of these causing the game to end and payoffs to be distributed.*

This definition is quite a mouthful, but it formally captures the "physical" structure of a game tree, ignoring the actions of players and what they know when they move. To illustrate the definition, consider the following game tree with five players depicted in figure X.X:

[2] I borrow heavily from the notation in Fudenberg and Tirole (1991). This will serve the reader who is interested in referencing their more advanced treatment of the subject, and hopefully not deter the reader who is not.

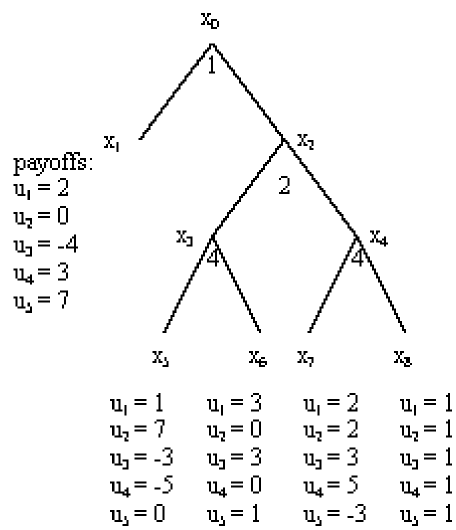

As the definition explains, every node (point in the game) can be reached as a consequence of actions that were chosen at the node that precedes it. The *root* is the beginning of the game, and here the node x_0 is the root at which the game begins with player 1 having two choices. The *terminal nodes* are the nodes at which the game can end, causing payoffs to be distributed. Here the set of terminal nodes is $Z = \{x_5, x_6, x_7, x_8\}$.

The payoffs for players $i \in N$ are given over terminal nodes: $u_i : Z \to \Re$, where $u_i(z)$ is i's payoff if terminal node z is reached. For example, if node x_5 is reached, then player 2 gets $u_2(x_5) = 7$ and player 5 gets $u_5(x_5) = 0$.

The precedence relation is derived from the actions of players as the game proceeds. For example, player 1's action at the root x_0 determines whether the game will terminate at node x_1 with payoffs $(2, 0, -4, 3, 7)$, or whether player 2 will get to play at node x_2. Player 2 then can choose whether player 4 will play at x_3 or at x_4, and at each of these player 4 has two choices that all end in termination of the game. Players 3 and 5 are not interesting since they have no moves to make.[3]

[3] This could easily be captured in normal forms as well by giving such "dummy" players a pure strategy set S_i that is a singleton, so that they have no choice to make.

Notice, however, that in figure X.X we have a particular order of play: first player 1, then 2, and then 4. How did we assign players to nodes, when this was not part of the definition of a game tree? Indeed, as noted earlier, the definition of a game tree is not complete enough to represent an extensive form game, and the following adds the way in which players are "injected" into the tree:

Definition 19 *The order of players is given by a function from non-terminal nodes to the set of players, $i : X \backslash Z \to N$, which identifies a player $i(x)$ for each $x \in X \backslash Z$.*[4] *The set of actions that are possible at node x are denoted by $A(x)$.*

Again, this formal definition may at first seem a bit abstract and detached, but it formally represents a simple idea. For example, in the game of figure X.X, $i(x_0) = 1$, which means that the player who moves at the root x_0 is player 1, and similarly $i(x_2) = 2$, and $i(x_3) = i(x_4) = 4$. As for what players can choose at each node, there are two elements of choice in each of the sets $A(x_0)$, $A(x_2)$, $A(x_3)$ and $A(x_4)$. If, for example, we denote the moves of player 1 at the root as "left" and "right", then we can write $A(x_0) = \{L, R\}$.[5]

There is still one missing component: how do we describe the knowledge of each player when he moves? In our example above, we see that player 2 moves after player 1, but does he know what player 1 did? For this example we would think the answer is obvious: if player 1 chose his other action, the game would end and player 2 would not have the option of moving. Looking at player 4, however, can offer something to dwell on. It seems implicit in the way we drew the game tree in figure X.X, that player 4 knows what happened before he moves, that is, if he is at x_3 or at x_4.

It might be, however, that player 4 needs to make his move *without knowing* what player 2 did at x_2. For this to be a conceivable situation, we need to find a way to represent the case in which player 4 *cannot distinguish* between himself being at x_3 or at x_4. That is, we need to make statements like "I know that I am at x_3 or at x_4, but I don't know at which of the two I am."

[4] The notation $x \in X \backslash Z$ is set-theoretic, and means "the elements x that are in X but not in Z".

[5] Note that the moves from any non-terminal node result in a move to another node in the game. Thus, we can save on notation and give moves "names" that are consistent with the nodes they will result in. For our current example we could write $A(x_0) = \{x_1, x_2\}$, $A(x_2) = \{x_3, x_4\}$, $A(x_3) = \{x_5, x_6\}$, and $A(x_4) = \{x_7, x_8\}$.

This argument implies that we need to put some formal structure on *the information* that a player has when it is his turn to move. A player can have very fine information and know exactly where he is at, or he may have coarser information and not know what has been done before his move, therefore not knowing exactly where he is in the game tree. For this we introduce the useful notion of an information set:

Definition 20 *Every node x has an* information set $h(x)$ *that partitions the nodes of the game. If $x \neq x'$ and if $x' \in h(x)$, then the player who moves at x does not know whether he is at x or x'.*

Once again, the formal definition has a simple idea behind it. Imagine some node x at which a player has to move, for example, x_3 at which player 4 has to move in the game above. We know now that player 4 is associated with a move at x_3, and now we want to describe whether or not he knows that he is at x_3. We can think of the information set $h(x_3)$ as *what player 4 knows at* x_3. If $h(x_3) = \{x_3\}$, this means that the information set at x_3 includes only the node x_3 (in which case, $h(x_3)$ is a *singleton* because it includes only one element). This means that player 4 has information that says "I am at x_3".

If, in contrast, we want to represent a game in which player 4 does not know whether he is at x_3 or x_4, then it must be the case that his information is "I am either at x_3 or x_4, but I don't know at which of these two." Thus, we will write $h(x_3) = \{x_3, x_4\}$, which exactly means that when player 4 is at x_3, he does not know whether he is at x_3 or x_4. Note that there is an immediate consequence of this representation of knowledge: if $x_4 \in h(x_3)$ then $x_3 \in h(x_4)$. This must be true because if, when player 4 is at x_3, he cannot distinguish between x_3 and x_4, then it must be that he is unable to distinguish between the two if he is at x_4. Thus, more generally, if $x \neq x'$ and $x' \in h(x)$ then $x \in h(x')$.

How can we use the graphical representation of the tree to distinguish whether a player knows where he is or not? To do this, we draw "enclosures" to denote information sets. For example, in the game in figure X.X player 2 cannot distinguish between x_2 and x_3, so that $h(x_2) = h(x_3) = \{x_2, x_3\}$ and both nodes are encircled together to denote this information for player 2.

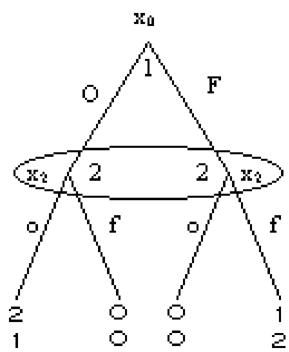

FIGURE 11.2.

A close look at figure X.X will make the game seem quite familiar. Indeed, this is the extensive form representation of the battle of the sexes that we have already seen and analyzed previously. Player 1 chooses from $S_1 = \{O, F\}$, and player 2 chooses between $S_2 = \{o, f\}$, without observing the choice of player 1.

The fact that a player cannot distinguish between two nodes that are in the same information set must immediately imply the following:

Fact: If $h(x)$ is not a singleton (single node), then

1. all $x' \in h(x)$ belong to the same player
2. if $x' \in h(x)$ then $A(x') = A(x)$

This fact should be quite obvious, but is worth mentioning. If, contrary to the first point, two nodes in the same information set belonged to two different players, it is as if neither player is sure if it is his move or not, a situation we deem as unreasonable. Furthermore, if contrary to the second point $x' \in h(x)$ but $A(x') \neq A(x)$, then by the mere fact that the player has different actions to choose from at each of the nodes x and x', he should be able to distinguish between them.

11.3 Imperfect versus Perfect Information

Recall that we defined *complete information* previously in part I as the situation in which each player i knows the payoff function of each $j \in N$, and this is common knowledge. This definition sufficed for the normal form representation. For extensive form games, however, it is useful to distinguish between two different types of complete information games:

Definition 21 *A game in which every information set $h(x)$ is a singleton is called a* Game of Perfect Information. *A game in which some information sets contain several nodes is called a* Game of Imperfect Information.

That is, in a game of perfect information every player knows exactly where he is in the game, while in a game of (complete but) imperfect information some players do not know where they are. Notice, therefore, that simultaneous move games fall into this second category because some information sets must include several nodes, like in the battle of the sexes shown above.

Games of imperfect information are also useful to capture uncertainty a player may have about acts of nature.

12

Strategies in the Extensive Form

In the normal form game it was quite easy to define a strategy for a player: a pure strategy was some element from his set of actions, S_i, and a mixed strategy was some probability distribution over these actions. It is very easy to extend this idea to extensive form games as follows:

Definition 22 *A* ***Pure Strategy*** *for player* i *is a* ***complete plan of play*** *saying what player* i *will do at each of his information sets.*

Take, for example, the battle of the sexes in figure X.X above. The pure strategies for player 1 are $S_1 = \{O, F\}$, and for player 2 are $S_2 = \{o, f\}$. This is a very simple example since each player has only one information set, making it identical to the simple normal form game we have already encountered.

Things can be just a bit trickier for the game depicted in figure X.X. In this game, which is the battle of the sexes but with player 2 observing the action that player 1 played, a "complete plan of play" must say what player 2 will do for *each* choice of player 1. That is, player 2's choice of action from $\{o, f\}$ can be made *contingent* on what player 1 will do. For this example, we can capture the set of pure strategies for player 2 as follows:

$$S_2 = \{oo, of, fo, ff\}$$

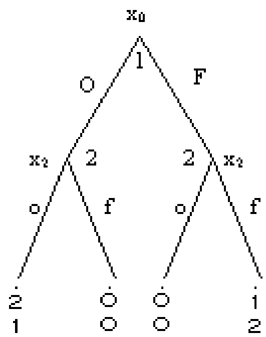

FIGURE 12.1.

where "ab" means "I will play a if player 1 plays O and I will play b if he plays F." For player 1 the set remains $S_1 = \{O, F\}$.

The example of figure X.X is illuminating. Even though player 2 has only two actions to choose from, by moving *after observing* what player 1 has chosen, he has *conditional strategies*. In this example the two actions translate into four pure strategies, due to the two information sets he has.

This means that a potentially small set of moves will translate into a much larger set of strategies when sequential moves are possible, and when players have knowledge of what preceded their play. In general, if a player i has $k > 1$ information sets, the first with t_1 actions to choose from, the second with t_2, and so on until t_k, then his total number of pure strategies will be,

$$|S| = t_1 \times t_2 \times \cdots \times t_k \,.$$

For example, a player with 3 information sets, 2 actions in the first, 3 in the second and 4 in the third will have a total of 24 pure strategies!

Once we have figured out what pure strategies are, the definition of mixed strategies follows immediately:

Definition 23 *A **Mixed Strategy** for player i is a probability distribution over his pure strategies.*

How do we interpret a mixed strategy? Exactly in the same way that it applies to the normal form: a player randomly chooses between all his complete plans of play, and once a particular plan is selected the player follows it. However, notice that this interpretation takes away some of the dynamic flavor that we set out to capture with extensive form games. Namely, when a mixed strategy is used, the player selects a plan and then follows a *particular pure strategy.*

This description of mixed strategies was sensible for normal form games since there it was a once-and-for-all choice to be made. In a game tree, however, the player may want to randomize at some nodes, independently of what he did in earlier nodes where he played. This cannot be captured by mixed strategies as defined above because one the randomization is over with, the player is choosing a pure plan of action.

To capture the possibility of players randomizing as the game unfolds, we introduce some notation that builds on what we have already developed. Let h_i represent some information set at which player i plays, and let H_i be the set of all information sets at which player i plays. Let $A_i(h_i)$ be the actions that player i can take at h_i, and let A_i be the set of all actions of player i, $A_i = \cup_{h_i \in H_i} A_i(h_i)$ (that is the union of all the elements in all the $A_i(h_i)$'s). We can now redefine a pure strategy, and introduce the new idea of a behavioral strategy:

Definition 24 *A **Pure Strategy** for player i is a mapping $s_i : H_i \to A_i$ that assigns an action $s_i(h_i)$ for all information sets $h_i \in H_i$. A **Behavioral Strategy** specifies for each information set $h_i \in H_i$, an independent probability distribution over $A_i(h_i)$, and is denoted by $\sigma_i : H_i \to \triangle A_i(h_i)$.*

One can argue, that a behavioral strategy is more loyal to the dynamic nature of the extensive form game. When using such a strategy, a player mixes between his actions whenever he is called to play. This differs from a mixed strategy, in which a player mixes before playing the game, but then is loyal to the selected pure strategy.

Luce and Raiffa (1957) provide a nice analogy for the different strategy types introduced above. A pure strategy can be thought of as an instruction manual in which each page tells the player which pure action to take at a particular information set, and the number of pages is the number of information sets the player has. The set S_i of pure strategies can, therefore, be treated like a library of such pure

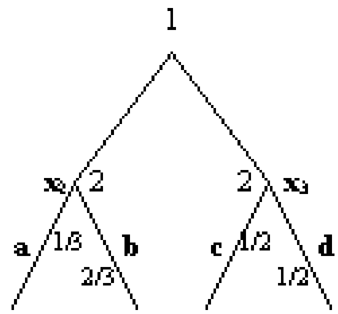

FIGURE 12.2.

strategy manuals. A mixed strategy is therefore choosing one of these manuals at random, and then following it.

In contrast, a behavioral strategy is a manual that prescribes possibly random actions on each of the pages associated with play at particular information sets. Consider the example in figure X.X.In it, player 2 has two information sets associated with the nodes x_2 and x_3, and in each he can choose between two actions. A pure strategy would be an element from $S_2 = \{ac, ad, bc, bd\}$. A mixed strategy would be a probability distribution $\{(p_{ac}, p_{ad}, p_{bc}, p_{bd}) : p_s \geq 0, \sum_{s \in S_2} p_s = 1\}$. A behavioral strategy for this example would be four probabilities, $\sigma_2(a), \sigma_2(b), \sigma_2(c)$, and $\sigma_2(d)$, where $\sigma_2(a)+\sigma_2(b) = \sigma_2(c)+\sigma_2(d) = 1$.[1] In the figure we have $\sigma_2(a) = \frac{1}{3}, \sigma_2(b) = \frac{2}{3}$, and $\sigma_2(c) = \sigma_2(d) = \frac{1}{2}$.

12.1 Mixed versus Behavioral Strategies

An interesting question is whether we are losing anything by restricting attention to either behavioral or mixed strategies as forms of randomized play. For the game in figure 3.5, for example, the set of mixed (not behavioral) strategies are probability distributions of the form $(p_{ac}, p_{ad}, p_{bc}, p_{bd})$, with $\sum_{s \in S_2} p_s = 1$. The first part of our question is that given a mixed strategy, can we find a behavioral strategy that leads to the same outcomes? The answer is yes, which is quite easy to see. Conditional

[1] Clearly, it would suffice to just give values for $\sigma_2(a)$ and $\sigma_2(c)$, and the other values would follow from the definition of a probability distribution adding up to one.

upon reaching x_2, the probability of playing a is $\Pr\{a|x_2\} = p_{ac} + p_{ad}$, and the probability of playing b is $\Pr\{b|x_2\} = p_{bc} + p_{bd}$, so that $\Pr\{a|x_2\} + \Pr\{b|x_2\} = 1$. Similarly, conditional upon reaching x_3, the probability of playing c is $\Pr\{c|x_3\} = p_{ac} + p_{bc}$, and the probability of playing d is $\Pr\{d|x_3\} = p_{ad} + p_{bd}$, so that $\Pr\{c|x_3\} + \Pr\{d|x_3\} = 1$. Thus, *any* mixed strategy can be achieved as a behavioral strategy.

The complementary question is whether we can find a mixed strategy for player 2 in figure 3.5 that would replicate the behavioral strategy $\sigma_2(a) = \frac{1}{3}, \sigma_2(c) = \frac{1}{2}$? To answer this consider the mixed strategy $(p_{ac}, p_{ad}, p_{bc}, p_{bd})$. Notice that *conditional* on player 1 choosing the node x_2, action a will be chosen with probability $p_{ac} + p_{ad}$, and action b will be chosen with probability $p_{bc} + p_{bd}$. Similarly, *conditional* on player 1 choosing the node x_3, action c will be chosen with probability $p_{ac} + p_{bc}$, and action d will be chosen with probability $p_{ad} + p_{bd}$. To replicate the behavioral strategy $\sigma_2(a) = \frac{1}{3}, \sigma_2(c) = \frac{1}{2}$, four equalities must be satisfied:

$$\begin{aligned}
\Pr\{a|x_2\} &= p_{ac} + p_{ad} = \frac{1}{3} \\
\Pr\{b|x_2\} &= p_{bc} + p_{bd} = \frac{2}{3} \\
\Pr\{c|x_3\} &= p_{ac} + p_{bc} = \frac{1}{2} \\
\Pr\{d|x_3\} &= p_{ad} + p_{bd} = \frac{1}{2}
\end{aligned}$$

It turns out that this system of equations is "unidentified" in the sense that many values of $(p_{ac}, p_{ad}, p_{bc}, p_{bd})$ will satisfy the four equations, meaning that this particular behavioral strategy can be generated by many different mixed strategies. For example, $(p_{ac}, p_{ad}, p_{bc}, p_{bd}) = (\frac{1}{3}, 0, \frac{1}{6}, \frac{1}{2})$ will lead to equivalent outcomes to the behavioral strategy $\sigma_2(a) = \frac{1}{3}, \sigma_2(c) = \frac{1}{2}$. Another mixed strategy that does the trick is $(p_{ac}, p_{ad}, p_{bc}, p_{bd}) = (\frac{1}{6}, \frac{1}{6}, \frac{1}{3}, \frac{1}{3})$, and many more (in fact, a continuum) will work.

This example suggests that any mixture that can be represented by a mixed strategy can also be represented by a behavioral strategy, and vice versa. Indeed, this result is true if players never "forget" what they had previously done. Formally:

Definition 25 *A game of* Perfect recall *is one in which no player ever forgets information that he previously knew.*

For the class of perfect recall games, Kuhn (1953) proved that mixed and behavioral strategies are equivalent, in the sense that given strategies of i's opponents, the same distribution over outcomes can be generated by either a mixed or a behavioral strategy of player i.

Though we will only be concerned with games of perfect recall, it is useful to see a simple example of games without perfect recall.

Example: The Absent Minded Driver

TBA

12.2 Normal-Form Representation of Extensive-Form Games

Consider the two variants of the battle of the sexes given above in figures 3.3 and 3.4. The first is equivalent to the original game we analyzed in normal form, and indeed, can be translated immediately into the normal form as follows:

	o	f
O	2,1	0,0
F	0,0	1,2

As it turns out, *any extensive form game* can be turned into a normal form game that is given by the set of players, the set of derived pure strategies, and the payoffs resulting from the actual play of any specified profile of strategies.

If there are two players and finite strategy sets, the game can be represented by a matrix as we have discussed in the previous section on normal form games. Take the sequential move battle of the sexes game depicted in figure 3.4. Recall that $S_1 = \{O, F\}$ and that $S_2 = \{oo, of, fo, ff\}$ where of, for example, means that player 2 plays o after player 1 plays O while player 2 plays f after player 1 plays F. This game can be represented by a 2×4 matrix as follows:

	oo	of	fo	ff
O	$2,1$	$2,1$	$0,0$	$0,0$
F	$0,0$	$1,2$	$0,0$	$1,2$

As this matrix demonstrates, each of the four payoffs in the original extensive form game is replicated twice. This happens because for a certain choice of player 1, two pure strategies of player 2 are equivalent. For example, if player 1 plays O, then only the first component of player 2's strategy matters, so that oo (player 2 playing o after O and o after F) and of (player 2 playing o after O and f after F) yield the same outcome, as do fo and ff. If, however, player 1 plays F, then only the second component of player 2's strategy matters, so that oo and fo yield the same outcome, as do of and ff.

Clearly, this exercise of transforming extensive form games into the normal form misses the dynamic nature of the extensive form game. Why, then, would we be interested in this exercise? It turns out to be very useful to find the Nash equilibria of the original extensive form game, especially when we have the ability to write the normal form in a matrix.

12.3 Nash Equilibrium and Paths of Play

Now that we have completed the description of dynamic games as extensive form games, we can move on to analyze the reasonable predictions using the concept of Nash equilibrium. Consider again the standard battle of the sexes in figure 3.3, by its equivalence to the normal form that we had already analyzed, we know that there are two pure strategy Nash equilibria, (O, o) and (F, f), and a mixed strategy one that we solved for earlier.

For the game depicted in figure 3.4, these equilibria can be replicated by (O, oo) and (F, ff). Are there other Nash equilibria in pure strategies? A glance at the extensive form does not immediately answer this. Now look again at the normal form,

	oo	of	fo	ff
O	$\overline{\underline{2,1}}$	$\overline{\underline{2,1}}$	$\underline{0,0}$	$0,0$
F	$0,0$	$\overline{1,2}$	$\underline{0,0}$	$\overline{\underline{1,2}}$

Sequential Battle of the Sexes

and perform the under/over-line method to reveal a third pure strategy Nash equilibrium: (O, of), which yields the same outcome as (O, oo).

This is a convenient feature of the normal form representation of an extensive form game: it will immediately reveal all the pure strategies of each player, and in turn will lead us to easily see the pure strategy profiles that are Nash equilibria. (That is, easy in matrix games, and will requite a bit more work for other games by finding the best response correspondences.)

In the extensive form, any Nash equilibrium is not only a prediction of the outcomes through the terminal nodes, but also a prediction about the *path of play,* also called the *equilibrium path* . In a Nash equilibrium, players choose to proceed on the equilibrium path because of their beliefs about what the other players are doing, which has consequences on the outcomes realized when a player moves *off the equilibrium path*, i.e., deviates from his proposed strategy.

Another way of putting this is that in a Nash equilibrium, players choose to proceed with the equilibrium path because of their beliefs about what happens *on and off the equilibrium path.* Consider again the sequential battle of the sexes:

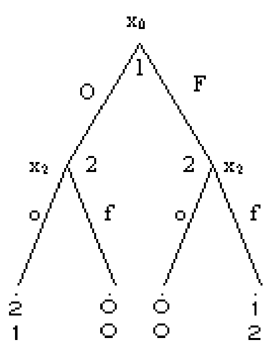

In this game, we saw through the conversion to a normal form matrix, that one equilibrium profile is (F, ff), resulting in payoffs of $(1, 2)$. This equilibrium is supported by the fact that player 1 *believes* that if he would deviate and play O, then he would receive 0 because player 2 will proceed to play f. That is, player 2's "threat" of how he will proceed *off the equilibrium path* is supporting the actions of player 1 *on the equilibrium path.* L8

13

Credibility and Sequential Rationality

Consider again the extensive and normal forms of the sequential (perfect information) Battle of the Sexes. Now ask yourself: is the Nash equilibrium (F, ff) a reasonable prediction about the rational choice of player 1? By the definition of Nash equilibrium it is: ff is a best response to F, and vice versa. However, this implies that if, for some unexpected reason, player 1 would suddenly choose to play O, then player 2 would *not respond optimally*: her strategy ff commits her to choose f even though o would yield her a higher utility after a choice of O by player 1.

This sheds light on a weakness of the normal form. The normal form representation treats all choices as "once-and-for-all" simultaneous choices, and thus beliefs can never be challenged. We would, however, expect rational players to play optimally in response to their beliefs *whenever* they are called to move. The normal form representation of an actual sequential game is not very meaningful to address such a requirement of *sequential rationality*. This requirement will put more constraints on what we would tolerate as "rational behavior", since we should expect players to be *sequentially rational*, which is the focus of this chapter.

This reasoning would suggest that of the three Nash equilibria in this game, two seem somewhat unappealing. Namely, the equilibria (O, oo) and (F, ff) have player 2 committing to a strategy that, despite being a best response to player 1's strategy, would not have been optimal were player 1 to deviate from his strategy and cause the game to move *off the equilibrium path*. In what follows, we will set up some structure that will result in more refined predictions for dynamic games. These will indeed rule out such equilibria, and as we will clearly see later, will only admit the equilibrium (O, of) as the unique equilibria that survives the more stringent requirements.

13.1 Sequential Rationality and Backward Induction

To address the critique that we posed in the previous example about the equilibria (O, oo) and (F, ff), we will directly criticize the behavior of player 2 in the event that player 1 did not follow his prescribed strategy. That is, we will insist that a player use strategies that are optimal at *every node in the game tree*. We call this principle *sequential rationality*, since it implies that players are playing rationally at every stage in the sequence of the game, regardless of it being on or off the equilibrium path of play.

Going back to the game above, we ask: what should player 2 do in each of her information sets? The answer is obvious, if player 1 played O, then player 2 should play o, and if player 1 played F, then player 2 should play f. Any other strategy should not be played in each of these information sets, which implies that player 2 should be playing the pure strategy of.

Now move back to the root of the game where player 1 has to choose between O and F. Taking into account the sequential rationality of player 2, player 1 should conclude that playing O will result in the payoffs $(2, 1)$ while playing F will result in the payoffs $(1, 2)$. Now applying sequential rationality to player 1 implies that player 1, who is correctly predicting the behavior of player 2, should choose O, and the unique prediction from this process is the path of play O followed by o. Furthermore, the process predicts what would happen if players *deviate from the path of play*: if player 1 chooses F then 2 will choose f. We conclude that the Nash equilibrium (O, of) uniquely survives this procedure.

This type of procedure, which starts at nodes that precede only terminal nodes at the end of the game and moves backward, is known as *backward induction*. It turns out, as the example above suggests, that when we apply this procedure to *finite games of perfect information*, then we will get a specification of strategies for each player that are sequentially rational. By finite we mean that the game has a finite number of sequences, after which it ends.

In fact, by the construction of the backward induction procedure, each player will play a best response to the other players' actions, which results in the following theorem:

Theorem 9 *Backward induction can be applied to any finite game of perfect information, and will result in a sequentially rational Nash equilibrium. Furthermore, if no two terminal nodes prescribe the same payoffs to any player, this procedure will result in a unique sequentially rational Nash equilibrium.*

This result is known as *Zermelo's Theorem,* even though it was proven in 1913, before John Nash was even born. Zermelo's original theorem stated that a game like chess, which has a finite number of moves at each stage and that must end in a finite number of stages, must have a solution. We can draw up a game tree, and apply this process to find the sequentially rational plays at each stage of the game, dependent on any history, taking into account the future stages that have already been solved for.

The proof is rather straightforward, and in based on a simple "backward-induction" argument, or multi-person dynamic program based on Zermelo's algorithm. Since it is a finite game of perfect information, we can identify all the terminal nodes, and the players who are at nodes that immediately precede a terminal node. Let every such "level 1" player at a node once-removed from a terminal node choose the action that maximizes his payoff.[1] Now similarly define "level 2" players, and let them choose the action that maximizes their payoff *given* that level 1 players will choose their action as specified before. This process continues iteratively until

[1] This is where the qualifier "if no two terminal nodes prescribe the same payoffs to any player" plays a role. If it holds, then there is a unique action that maximizes each player's payoff. If not, we can select any one of the maximizing actions.

we reach the root of the game, and results in a specification of strategies that are a Nash equilibrium and are sequentially rational by construction.

It is worth noting that this theorem implies that every finite game of perfect information has at least one pure strategy Nash equilibrium, which is the sequentially rational one. As the example we dealt with above implies, other pure strategy Nash equilibria may exist.

To solidify the procedure of backward induction, consider the game tree in figure X.X below:

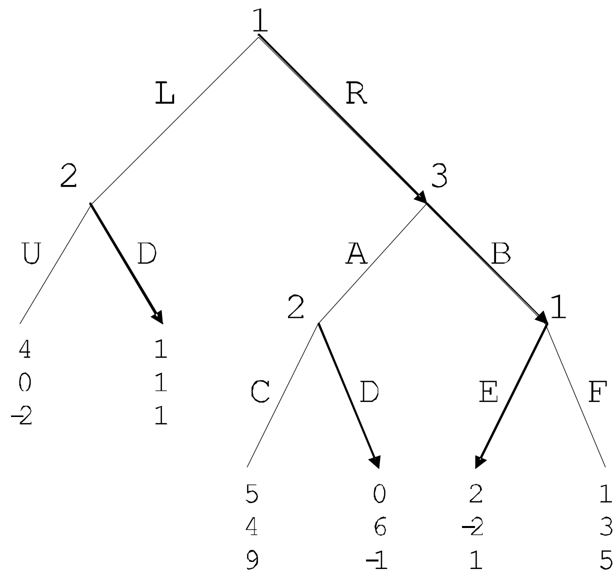

There are three nodes that precede the terminal nodes, two for player 2 and one for player 1. The dark arrows determine what every player would do at his information set, and we can continue backward to determine the sequentially rational play of players at "higher up" preceding nodes, assuming that the players that follow after them will adhere to their sequential rational play. At the nodes that precede only terminal nodes, player 2 plays D at each of his information sets, and player 1 plays E. Player 3 can choose between A, which through player 2's action yields player 3 a payoff of (-1), or he can choose B, which through player 1's action yields player 3 a payoff of 1. Thus, player 3 will choose B. Now player 1 can choose between a payoff of 1 (L followed by D), or a payoff of 2 (R followed by B and then E). Thus, player 1 will choose R. We can write this down as a pure strategy Nash equilibrium: (RE, DD, B) which are the pure strategies that result from this backward induction procedure. (I am skipping the obvious part of defining what a strategy here is for each player, but this must be made clear for the sake of precision).

As a final example consider the Stackelberg model of duopoly that we introduced at the beginning of this chapter. The best response of firm 2 was

$$q_2 = \frac{100 - q_1}{2}\,, \tag{13.1}$$

and sequential rationality implies that player 2 should always choose q_2 accordingly. That is, given *any quantity* that firm 1 chooses, firm 2 *must play* according to its best reply in (13.1). This, in turn, means that a sequentially rational firm 1 will chooses q_1 to solve:

$$\max_{q_1} \left[100 - q_1 - \left(\frac{100 - q_1}{2}\right)\right] q_1$$

that results in the Stackelberg outcome of $q_1 = 50$, $q_2 = 25$.

There is a simple, yet *very important* point worth emphasizing here. When we write the pair of strategies down as a pure strategy Nash equilibrium *we must be careful* to specify the strategies correctly: player 2 has a continuum of information sets, each being a particular choice of q_1. This implies that the backward induction solution yields the following Nash equilibrium: $(q_1, q_2(q_1)) = \left(50, \frac{100-q_1}{2}\right)$. Writing down $(q_1, q_2) = (50, 25)$ is not a Nash equilibrium because $q_1 = 50$ is not a best response to $q_2 = 25$ (check this!)

13.2 Subgame Perfect Nash Equilibrium

We saw that Zermelo's theorem is useful in helping us identify sequentially rational Nash equilibria for a large class of games. However, this is a *solution procedure*, not a *solution concept* in that it is well defined for some games but not for others. The reason is that this procedure only applies to finite games of perfect information, whereas a solution concept such as Nash equilibrium applies to *all* extensive form games (by the fact that they can be represented as a normal form game).

Our next goal is to find a natural way to extend the ideas of sequential rationality to games that are not constrained to be finite, or of perfect information. Intuitively, when a game id not a game of perfect information then a player has non-degenerate information sets, which in turn implies that we may not be able to identify his best action. The reason is that his best action may depend on which node he is at within his information set, which itself depends on the actions of players that moved before him. Thus, we will have to consider a solution concept that looks at sequential rationality of the game with these type of dependencies. For this reason we define:

Definition 26 *A proper subgame, G, of an extensive-form game Γ consists of a single node and all its successors in Γ with the property that if $x' \in G$ and $x'' \in h(x')$, then $x'' \in G$. The subgame G is itself a* game tree *with its information sets and payoffs "inherited" from Γ.*

The idea of a proper subgame (which we will often just call a subgame) is simple, and gives us the ability to potentially "dissect" an extensive form game into a sequence of smaller games. However, we will require that every such smaller game is indeed an extensive form game, which means that it must have a unique root, and follow the structure that we have defined above. For example, consider the simultaneous move Battle of Sexes in figure X.X.

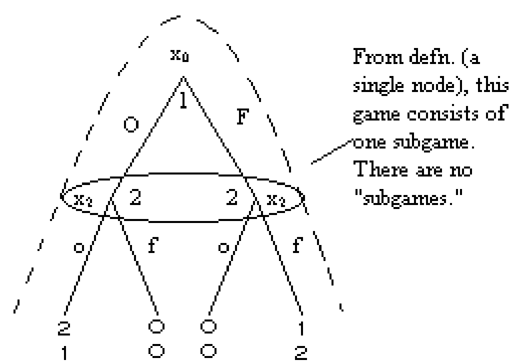

There is only one proper subgame in this game: the whole game. The reason is, by definition, that x_1 or x_2 cannot be roots of a subgame since they belong to the same information set, one that has multiple nodes.

Another example in which the only proper subgame is the whole game appears in figure X.X:

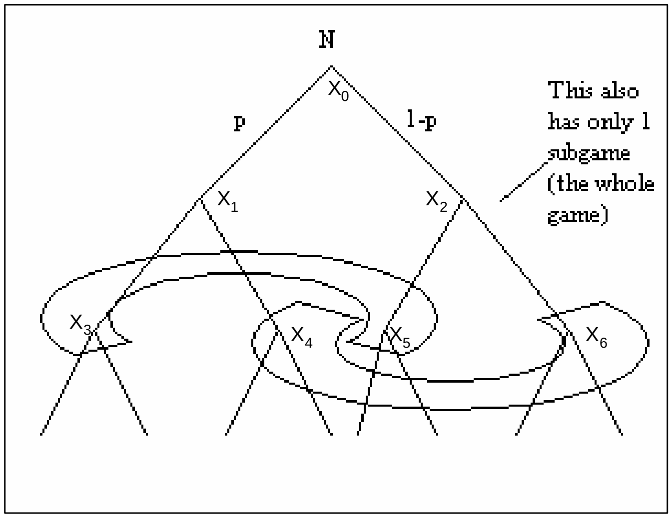

If, for example, one tries to have x_1 as the root of a proper subgame, then both x_3 and x_4 must be in that subgame. But then, because x_5 is in the same information as x_2, and x_6 is in the same information as x_3, both of these nodes must be included as well. This, however, does not result in a proper subgame, and a similar outcome would happen if we try to have x_2 as the root of a subgame. Thus, the only proper subgame is the whole game.

For the sequential battle of the sexes in figure X.X, there are three proper subgames. The whole game, and the two subgames with x_1 and x_2 as their roots. Similarly, you should be able to convince yourself that the three player game in figure X.X above has 5 proper subgames, each having a root that corresponds to one of the non-terminal nodes of the game. We are now ready to introduce the desired solution concept:

Definition 27 *A behavioral strategy profile σ of an extensive form game is a* ***subgame-perfect Nash equilibrium*** *if the restriction of σ to G is a Nash equilibrium for every proper subgame G of Γ.*

This important concept was introduced by Reinhard Selten (1975), who was the second of the three Nobel Laureates sharing the prize for game theory in

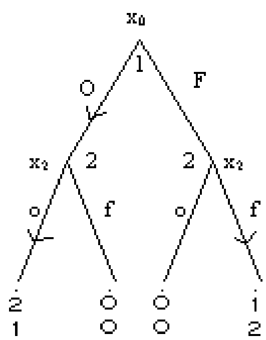

FIGURE 13.1.

1994. This equilibrium concept introduces sequential rationality into the concept of Nash equilibrium. It says that an equilibrium should not only be a profile of best responses *on the equilibrium path*, which is a necessary condition of a Nash equilibrium, but that the profile of strategies should be mutual best responses *off the equilibrium path as well.*

This logic follows from the requirement that the restriction of σ be a Nash equilibrium in every proper subgame. This stronger requirement of subgame-perfection requires the strategies that players are choosing to be mutual best responses even in those subgames that are not reached in equilibrium.

Notice that by the definition of a Subgame Perfect Nash equilibrium (SPE), every SPE is a Nash equilibrium. However, not all Nash equilibria are necessarily SPE, implying that SPE *refines* the set of Nash equilibria, yielding more refined predictions on behavior. To see this consider the sequential battle of the sexes in figure X.X, and it's corresponding normal form given as:

	oo	of	fo	ff
O	$\underline{\overline{2,1}}$	$\underline{\overline{2,1}}$	$\underline{0,0}$	$0,0$
F	$0,0$	$\overline{1,2}$	$\underline{0,0}$	$\underline{\overline{1,2}}$

in which we have identified three pure strategy Nash equilibria, (O, oo), (F, ff) and (O, of). Of these three, only (O, of) is the unique subgame-perfect equilibrium. This follows because in the subgame beginning at x_1, the only Nash equilibrium is player 2 choosing o, since she is the only player in that subgame and she must choose a best response to her belief, which must be "I am at x_1." Similarly, in the subgame beginning at x_2, the only Nash equilibrium is player 2 choosing f. Thus, of the three Nash equilibria of the whole game, only (O, of) satisfies the condition that its restriction is a Nash equilibrium for every proper subgame of the whole game.

To see the application of subgame-perfection in a game with continuous strategy sets, consider the Bertrand model that can be loosely illustrated as the game tree in figure X.X:

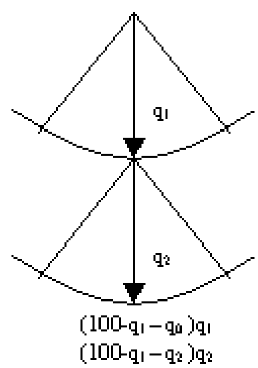

This describes the game with demand $p = 100 - q_1 - q_2$, and zero costs to the firms. Furthermore, firm 2 observes the quantity choice of firm , which we describe by having no information sets in the game tree. We can loosely depict the simultaneous Cournot game in figure X.X:

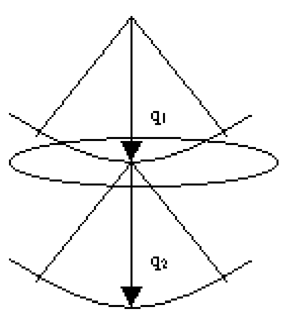

in which firm 2's information set includes all the choices of q_1 that firm 1 can make.

In the Cournot game we know from previous analysis that $q_1 = q_2 = 33\frac{1}{3}$ is the unique Nash equilibrium. Could this be part of a Nash equilibrium in the

Stackelberg game? The answer, maybe somewhat surprisingly is yes, and to see this consider the following strategies in the sequential game: $q_1 = 33\frac{1}{3}$, and $q_2(q_1) = 33\frac{1}{3}$ for any q_1. These two strategies are mutual best responses, implying that they constitute a Nash equilibrium. Notice two things: First, we have to specify q_2 for *every information set* for it to be a well defined strategy, and in this case we defined it to be $q_2(q_1) = 33\frac{1}{3}$ for every q_1. Second, it is a pair of mutual best responses *only on the equilibrium path.* That is, conditional on $q_1 = 33\frac{1}{3}$ then q_2 is a best response, and conditional on $q_2(q_1) = 33\frac{1}{3}$ for every q_1 then q_1 is a best response. This is all that Nash equilibrium requires.

However, as you might already have noticed, the "constant" strategy of player 2 is not sequentially rational because it is never a best response *off the equilibrium path* for $q_1 \neq 33\frac{1}{3}$. We know this from the backward induction argument presented earlier, after which we concluded that sequential rationality implies that

$$q_2 = \frac{100 - q_1}{2},$$

which in turn implied that $q_1 = 50$ and $q_2 = 25$ are the result of backward induction. Not surprisingly, the *unique* subgame-perfect equilibrium of the Stackelberg game is,

$$q_1^* = 50\ ; \ \ q_2^*(q_1) = \frac{100 - q_1}{2}\ .$$

Remark 7 *Interestingly, there are many other Nash equilibria that pop up when the game is sequential in the Stackelberg manner. To see this consider the following strategy for player 2:*

$$q_2'(q_1) = \begin{cases} 40 & \text{if } q_1 = 20 \\ 100 & \text{if } q_1 \neq 20 \end{cases} .$$

What would be the best response of firm 1? Clearly, $q_1 = 20$ yields positive profits, while any other choice yields non-positive profits. Thus, $q_1 = 20$ is the best response. But $q_2 = 40$ is a best response to $q_1 = 20$, implying that $q_1 = 20$ and $q_2'(q_1)$ constitute a Nash equilibrium. At this stage you should realize that there are many Nash equilibria, in fact, a continuum of Nash equilibria! To see this fix any $\widetilde{q}_1 \in [0, 100]$, and replace the 40 in $q_2'(q_1)$ with 2's best response $\widetilde{q}_2 = \frac{100 - \widetilde{q}_1}{2}$.

Notice that the two games we analyzed above were finite games of perfect information. For these games we had an easy way to find the subgame perfect equilibria, which was using the procedure of Backward induction. This is no coincidence:

Fact: *For any finite game of perfect information, the set of subgame-perfect Nash equilibria coincides with the set of Nash equilibria that survive backward induction.*

This is a very useful fact because it gives us an operational procedure to find all the SPE in finite games of perfect information. However, since this procedure does not apply to games of imperfect information, we need to consider a somewhat different procedure that mimics the workings of backward induction, as illustrated by the following example.

Example: Mutually Assured Destruction

This is a very nice example that appears in Gardner (2003, p.165). Two superpowers, players 1 and 2, have engaged in a provocative incident. The games starts with player 1's choice to either ignore the incident (I), resulting in the status quo with payoffs (0,0), or to escalate the situation (E). Following escalation by player 1, player 2 can back down (B), causing it to loose face and result in payoffs of $(1, -1)$ or it can choose to proceed to an atomic confrontation situation (A). Upon this choice, the players play a simultaneous move game where they can either retreat (R for player 1, r for player 2), or choose Doomsday (D for player 1, d for player 2) in which the world is destructed. If both call things off then they suffer a small loss due to the process and payoffs are $(-0.5, -0.5)$, and if either chooses doomsday then the world destructs and payoffs are $(-L, -L)$, where L is a very large number.

The extensive form of this game is depicted in figure X.X.

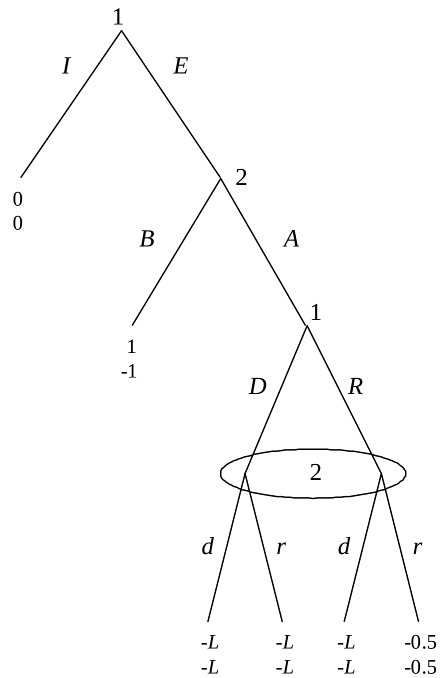

Before solving for the SPE, we begin by solving for the Nash equilibria of this game. To do this, it is convenient to transform the extensive form to a normal form game as follows:

	Br	Bd	Ar	Ad
IR	$\overline{0,0}$	$\overline{0,0}$	$\underline{\overline{\mathbf{0},\mathbf{0}}}$	$\underline{\overline{0,0}}$
ID	$\overline{0,0}$	$\overline{0,0}$	$\underline{\overline{0,0}}$	$\underline{\overline{0,0}}$
ER	$\underline{1,-1}$	$\underline{1,-1}$	$-\frac{1}{2},-\frac{1}{2}$	$-L,-L$
ED	$\underline{\overline{1,-1}}$	$\underline{\overline{\mathbf{1},-\mathbf{1}}}$	$-L,-L$	$-L,-L$

In this game, there are six pure-strategy Nash equilibria that are, (IR, Ar), (IR, Ad), (ID, Ar), (ID, Ad), (ED, Br) and (ED, Bd). (We will ignore the mixed strategy Nash equilibria in this example, and there will be many!)

How can we solve for the SPE? We use *backward induction over the subgames*. There are three proper subgames: the first is the whole game, the second is the point where 2 chooses between B and A, and the third starts where 1 chooses between R and D. (Notice that we can switch the order of play in the third subgame since it is a "simultaneous move subgame".)

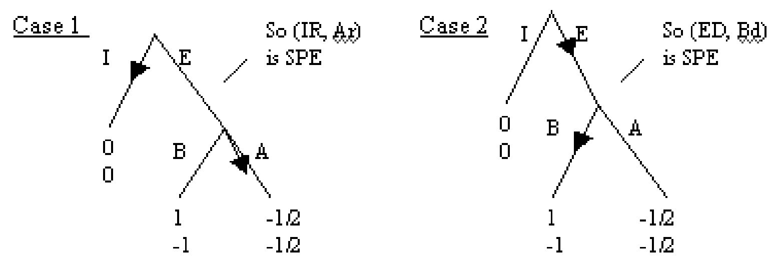

FIGURE 13.2.

It is easy to see that the third subgame is just the following normal form game,

$$\begin{array}{ccc} & r & d \\ R & -\frac{1}{2}, -\frac{1}{2} & -L, -L \\ D & -L, -L & -L, -L \end{array}$$

and this game has two pure strategy Nash equilibria:

- case 1: (R, r) that yields a payoff of $(-\frac{1}{2}, -\frac{1}{2})$
- case 2: (D, d) that yields a payoff of $(-L, -L)$

We can take the outcome of each of these cases, and "squeeze" the original extensive form game to one of the two cases described in figure 3.14.In this figure we see that for case 1, player 2 would prefer to choose A over B, and moving backward another step shows that player 1 would prefer I to E. On the other hand, for case 2 player 2 would prefer to choose B over A, and moving backward another step shows that player 1 would prefer E to I.

Thus, we have found two different subgame-perfect equilibria, (IR, Ar) and (ED, Bd). The reason we have two different subgame-perfect equilibria is because in the final subgame there are two different Nash equilibria. These two Nash equilibria have different implications on the behavior of player 2 in the preceding stage of the game, which in turn affects the behavior of player 1 at the beginning of the game.

There is a nice qualitative difference of the two equilibria. The first, (IR, Ar), in one in which player 1 chooses to ignore the incident because it believes that if it escalates then two things will happen: player 2 will go nuclear and they will then retreat together, resulting in a payoff of -0.5 instead of 0. On the other hand, (ED, Bd) represents the case that player 1 escalates because it believes that 2 will treat this as a signal that player 1 is willing to "go all the way", and in turn player 2 will back off. Gardner interprets this game as representing the Cuban missile crisis of 1962, in which the crisis started with the U.S. discovery of soviet nuclear missiles in Cuba, after which the U.S. escalated the crisis by quarantining Cuba. the U.S.S.R. then backed down, agreeing to remove its missiles from Cuba. We can view this as the U.S. having a credible threat: "if you don't back off we both die", which is a Nash equilibrium of the last stage of the game. L9