Unit 7.3: Sequential Games

Michael Malcolm

June 18, 2011

Sequential games are games in which players choose their strategies in some pre-determined order. The important difference with simultaneous games is that players who move later in the game can condition their choices on observed moves that were made earlier in the game. By contrast, the nature of simultaneous games implies that players must all choose their own strategies without knowing what strategies are chosen by other players.

1 Backwards Induction

Sequential games are solved via *backwards induction* – starting at the end of the game and working backwards. The equilibrium outcome derived through backwards induction is known as a *subgame perfect equilibrium.*

For a simple example, first consider the following simultaneous game.

	Y	**Z**
A	5,2	3,1
B	6,3	4,4

The Nash Equilibrium of the simultaneous game is (B,Z). Suppose instead that the game is played sequentially, with player 1 selecting his strategy first. Player 2 then observes player 1's choice and selects his own strategy. Note that, from a strategic perspective, this is now a completely different game since the information structure is different.

Sequential games are represented by game trees that are read from top to bottom, in order of the moves. The game tree for this new game is illustrated in figure 1.

Some terminology here. A *decision node* is where a player makes a choice. In this game tree, player 1 has one decision node and player 2 has two decision nodes. A *terminal node* occurs where no player has any other moves; in other words, where the game ends. Terminal nodes show the payoffs to the players for the corresponding outcome of the game when that terminal node is reached. Payoff numbers are written in the usual order, with player 1's payoff first, etc...

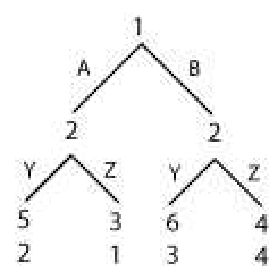

Figure 1: Game tree

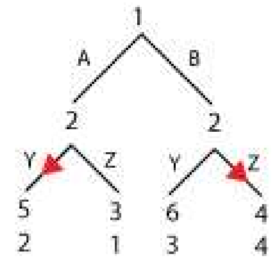

Figure 2: Game tree showing player 2's choices

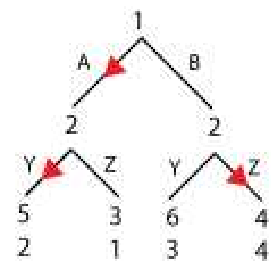

Figure 3: Game tree showing all players' choices. SPE outcome is (A,Y)

We solve sequential games by working backwards. When player 2 is at his left decision node, he gets a payoff of 2 by playing Y and a payoff of 1 by playing Z. Thus, at his left decision node, player 2 selects Y. On the other hand, when player 2 is at his right decision node, he gets a payoff of 3 by playing Y but a payoff of 4 by playing Z. At his right decision node, player 2 selects Z. Figure 2 shows arrows indicating player 2's decisions.

Working backwards, we can now determine player 1's optimal strategy. The point is that, when player 1 moves, he knows how player 2 is going to respond. This is why it is important to think backwards. If player 1 picks A, player 2 will respond by picking Y and so player 1's payoff would be 5. If player 1 picks B, player 2 will respond by picking Z and so player 1's payoff would be 4. Since he gets a payoff of 5 by selecting A and a payoff of 4 by selecting B, player 1 selects A initially. Figure 3 shows the choices of both players.

The subgame perfect equilibrium outcome of the game is for player 1 to select A and for player 2 to select Y.

2 Strategy Specification

There is a subtlety with specifying strategies in sequential games. A strategy in a sequential game needs to include directions for what the player will choose at *every* decision node, even decision nodes that are not reached. In the game depicted in figure 3, player 1 has two strategies: A and B. Player 2, however has four strategies, since we must specify a strategy *at each node*. Player 2's strategies are: YY, YZ, ZY, and ZZ, indicating her choices at the left node and the right node respectively.

So the equilibrium *outcome* in the game shown in figure 3 is that player 1 selects A and player 2 follows by selecting Y. However, the subgame perfect equilibrium is (A; YZ). The equilibrium must specify each player's complete strategy, which includes strategies chosen at every node, even nodes that are not reached.

In fact, a strategy in a sequential game needs to specify actions at decision nodes even when a player's *own* actions prevent the decision node from being reached. As an example, consider the game in figure 4.

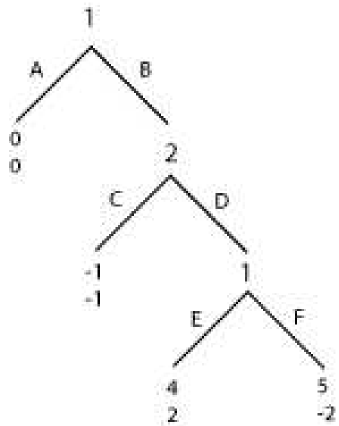

Figure 4: Game tree

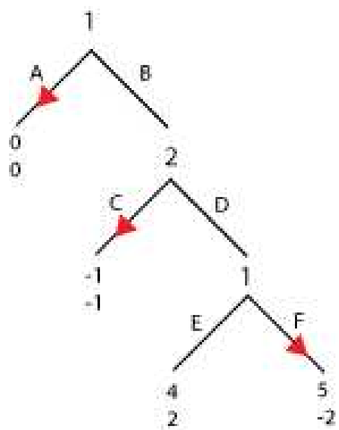

Figure 5: Game tree with decisions shown

Player 2 obviously has two strategies: C and D. Player 1 has *four* strategies: AE, AF, BE and BF, specifying one choice at each node. It seems odd that AE and AF are different strategies. After all, if player 1 selects A, then the game ends and so he never chooses between E and F. However, a strategy must specify a decision at every node.

Solving this game backwards: At his last decision node, player 1 would select F (4 from E versus 5 from F). This means that player 2 selects C (-2 from D since player 1 follows with F, but -1 from C). Moving backwards, player 1 selects A to begin with. These choices are shown with arrows in figure 5.

The *outcome* of this game is that player 1 selects A and the game is over. The subgame perfect equilibrium includes the whole strategy description: (AF, C). Why is this important? The reason that player 1 selects A in the first place is because player 2 would have selected C at his decision node. But the reason that player 2 selects C at his decision node is because player 1 would have selected F at his second decision node. After all, if player 1 would have selected E, then player 2 would select D and player 1 would select B initially. In this sense, it is important to be clear that player 1's strategy is AF rather than AE, even though the decision node between E and F is never actually reached.

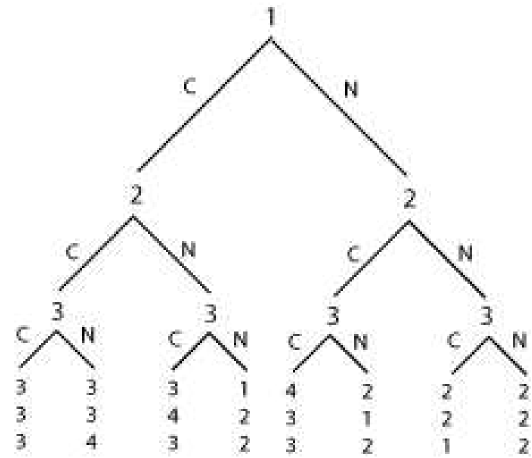

Figure 6: Game tree for public good provision game

3 Example: Contributions to a Public Good

Three players move sequentially and can choose whether to contribute (C) or not contribute (N) to the construction of a public good. If there are a total of 2 or 3 contributors, then the value of the public good to each person is 4. If there are a total of 0 or 1 contributors, then the value of the public good to each person is 2. Contributing costs 1.

The game tree is depicted in figure 6. Each player's payoff is 2 if there are zero or one contributors and 4 if there are two or three contributors, with a deduction of 1 if the player himself is a contributor.

Solving the game via backwards induction: Player 3's payoff-maximizing choices are N, C, C and N – on the decision nodes moving from left to right, respectively. Conditional on player 3's choices, player 2 chooses N at his left decision node (payoff of 3 from C and 4 from N, conditional on player 3's choices). Similarly, player 2 chooses C at his right decision node.

With this in mind, player 1 realizes that C will be followed by player 2 playing N and player 3 playing C – giving player 1 a payoff of 3. If player 1 plays N, this will be followed by player 2 playing C and player 3 playing C – giving player 1 a payoff of 4. Thus, player 1 selects N. The full backwards induction reasoning is shown in figure 7.

The equilibrium outcome is that player 1 does not contribute, while players 2 and 3 do. The subgame perfect equilibrium specifies the full strategy: (N, NC, NCCN). This describes player 1's, player 2's and player 3's strategies, respectively, at *all* their decision nodes, reading from left to right.

4 Normal Forms and Nash Equilibria of Sequential Games

Consider the game shown in figure 8.

Working backwards, player 1 at his bottom node chooses F (6 from F and 2 from E). Player 2 therefore chooses X at her left decision node (3 from W and 4 from X since player 1 follows with F) and chooses Y at her right decision node (5 from Y and 3 from Z). Moving backwards, player 1 begins by selecting U. The backwards induction logic is shown with arrows in figure 9.

Player 1 has four strategies in this game: UE, UF, DE and DF. Player 2 also has four strategies: WY, WZ, XY and XZ. Using these strategies, we can translate this game into the equivalent *normal form*, illustrated in the table below.

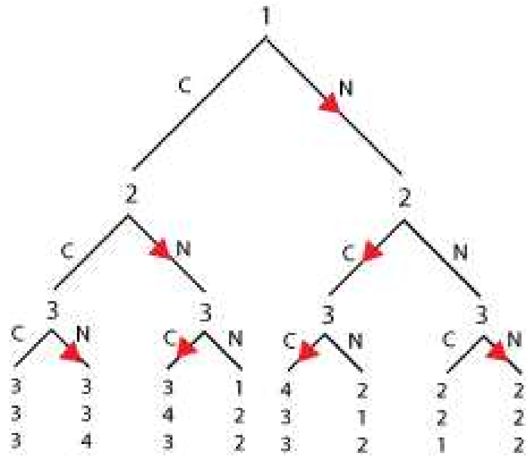

Figure 7: Game tree with backwards induction shown for the public good game

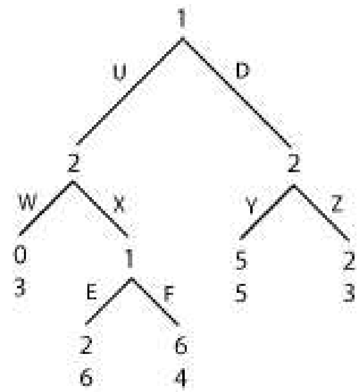

Figure 8: Sequential Game

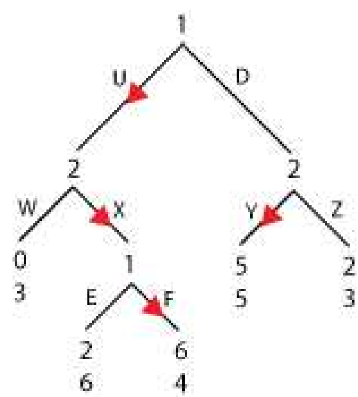

Figure 9: Sequential game with choices shown

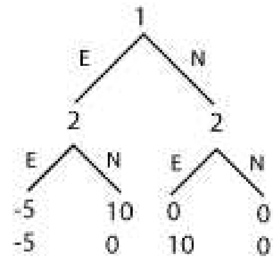

Figure 10: Entry deterrence game

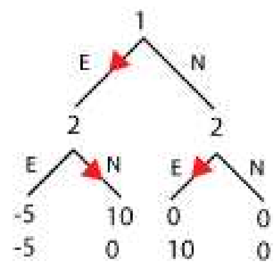

Figure 11: Choices in entry deterrence game

	WY	**WZ**	**XY**	**XZ**
UE	0,3	0,3	2,6	2,6
UF	0,3	0,3	6,4	6,4
DE	5,5	2,3	5,5	2,3
DF	5,5	2,3	5,5	2,3

To find the normal form of a sequential game, simply determine the payoffs at the terminal node when players follow the specified path. For example, when player 1's strategy is UE and player 2 plays WY, the players end up at the terminal node where the payoff is (0,3).

Standard best response analysis shows that this game has four Nash Equilibria: (UF,XY), (UF,XZ), (DE,WY) and (DE,WZ). However, looking back at figure 9, the subgame perfect equilibrium is (UF,XY). In general, the set of Nash Equilibria is larger than the set of subgame perfect equilibrium. A subgame perfect equilibrium is a Nash Equilibrium, but not vice versa – a Nash Equilibrium of a sequential game need not be a subgame perfect equilibrium.

5 Example: Entry Deterrence

Consider the following game, shown in figure 10. Player 1 can enter (E) or not enter (N) some market. After observing this choice, player 2 decides himself whether to enter or not. If neither enters, then both make 0. If both enter, the market is oversaturated and so both earn a loss of 5. However, if only one enters, then the entrant earns monopoly profit of 10.

Solving the game backwards: If player 1 enters, player 2 responds by not entering. If player 1 does not enter, player 2 responds by entering. Knowing this, player 1 enters. The SPE outcome is that player 1 enters and that player 2 does not. The full SPE is: (E,NE) – reflecting that player 2 does not enter at his left decision node, but would enter at his right decision node. The decisions are shown in figure 11.

However, notice something interesting. Suppose that player 2 could *commit* to entering regardless of whether player 1 enters. If player 1 actually believes this threat, then player 1 would not enter. After all, if player 1 believed that player 2 were going to always enter, then he would make -5 by entering and 0 by staying out.

The problem is that this threat seems unreasonable. If player 1 enters, is it really reasonable that player 2 is going to follow through on his threat and enter? Maybe – it depends on whether player 2 can make a *credible commitment.* This is a large area of research in economics, and the general idea of creating credible commitments applies to many settings other than firm entry deterrence.

WE can be more precise about this idea by considering the normal form. Player 1 has two strategies: E and N. Player 2 has four strategies: EE, EN, NE and NN, describing his possible choices at each decision node. The normal form is below.

	EE	**EN**	**NE**	**NN**
E	-5,-5	-5,-5	10,0	10,0
N	0,10	0,0	0,10	0,0

Recall that (E,NE) was the subgame perfect equilibrium of the sequential game. It is a Nash Equilibrium of the reduced normal form, as well.

Importantly, though (N,EE) is *also* a Nash Equilibrium of the reduced normal form. That is, if player 2 commits to EE, entering even when player 1 enters, then player 1's best response is to not enter. This is a Nash Equilibrium, but it is backed up by a threat that is not credible, and it is not a subgame perfect equilibrium of the sequential version of the game.

(E,NE) is a Nash Equilibrium and a subgame perfect equilibrium. (N,EE) is a Nash Equilibrium, but it is not a subgame perfect equilibrium. The distinction is that strategies in a Nash Equilibrium have to be rational only on the equilibrium path – players can make unreasonable threats off the equilibrium path. On the other hand, strategies in a subgame perfect equilibrium have to be rational both on and off the equilibrium path. Subgame perfect equilibrium is stricter than Nash Equilibrium.

6 Continuous Strategy Sets

Consider the game from the previous section where player 1 chooses a_1 and that player 2 chooses a_2. The player's payoffs are, respectively:

$$\Pi_1 = 3a_1 - a_1 a_2 - a_1^2$$
$$\Pi_2 = 4a_2 - a_1 a_2 - a_2^2$$

Suppose, instead of a simultaneous game, this is played as a sequential game with player 1 first selecting a_1, player 2 observing this choice, then selecting a_2. Solving the game backwards, we first determine player 2's best response to player 1's choice. Maximizing Π_2 with respect to a_2:

$$\frac{\partial \Pi_2}{\partial a_2} = 4 - a_1 - 2a_2 = 0$$
$$a_2 = \frac{4 - a_1}{2}$$
$$= 2 - \frac{1}{2}a_1$$

Player 1 now *anticipates* player 2's response to his choice of q_1. As a result, we can insert player 2's reaction function into player 1's payoff function.

$$
\begin{aligned}
\Pi_1 &= 3a_1 - a_1 a_2 - a_1^2 \\
&= 3a_1 - a_1\left(2 - \frac{1}{2}a_1\right) - a_1^2 \\
&= 3a_1 - 2a_1 + \frac{1}{2}a_1^2 - a_1^2 \\
&= a_1 - \frac{1}{2}a_1^2
\end{aligned}
$$

Player 1 now picks a_1 to maximize his payoff, having anticipated how player 2 will respond to his choices.

$$
\begin{aligned}
\frac{\partial \Pi_1}{\partial a_1} = 1 - a_1 &= 0 \\
a_1 &= 1
\end{aligned}
$$

Substituting back gives:

$$
a_2 = 2 - \frac{1}{2}a_1 = \frac{3}{2}
$$

So the SPE outcome of this sequential game is $(a_1, a_2) = \left(1, \frac{3}{2}\right)$. The SPE of the game specifies player 2's strategy conditional on any particular choice by player 1, so that would be expressed $(a_1, a_2) = \left(1, 2 - \frac{1}{2}a_1\right)$. The point is that we have specified player 2's strategy for any particular choice of a_1, not just for his particular choice in equilibrium $a_1 = 1$.

7 Stackelberg Oligopoly

Recall the Cournot oligopoly example from the previous section where two firms chose levels of output q_1 and q_2, respectively. Total market demand is $P = 140 - 2Q$, where Q is total market output $Q = q_1 + q_2$. The marginal cost of each unit of output was 20.

Recall that the profit function for firm 1 is:

$$
\begin{aligned}
\Pi_1 &= TR_1 - TC_1 \\
&= 120q_1 - 2q_1^2 - 2q_1 q_2
\end{aligned}
$$

and the profit function for firm 2 is:

$$
\begin{aligned}
\Pi_2 &= TR_2 - TC_2 \\
&= 120q_2 - 2q_2^2 - 2q_1 q_2
\end{aligned}
$$

Consider now the variant of this game where the firms choose their outputs sequentially rather than simultaneously. Firm 1 chooses q_1, then firm 2 observes this choice and chooses q_2. As usual, solving backwards, we first find the level of output that firm 2 will choose to maximize his profit:

$$
\begin{aligned}
\frac{\partial \Pi_2}{\partial q_2} = 120 - 4q_2 - 2q_1 &= 0 \\
4q_2 &= 120 - 2q_1 \\
q_2 &= 30 - \frac{1}{2}q_1
\end{aligned}
$$

Firm 1 anticipates this response and inserts firm 2's reaction function into his own profit function.

$$\begin{aligned}
\Pi_1 &= 120q_1 - 2q_1^2 - 2q_1q_2 \\
&= 120q_1 - 2q_1^2 - 2q_1\left(30 - \frac{1}{2}q_1\right) \\
&= 120q_1 - 2q_1^2 - 60q_1 + q_1^2 \\
&= 60q_1 - q_1^2
\end{aligned}$$

Firm 1 now chooses q_1 to maximize his profit:

$$\begin{aligned}
\frac{\partial \Pi_1}{\partial q_1} = 60 - 2q_1 &= 0 \\
q_1 &= 30
\end{aligned}$$

Substituting back:

$$q_2 = 30 - \frac{1}{2}q_1 = 15$$

So, in the Stackelberg equilibrium, firm 1 produces $q_1 = 30$, with firm 2 producing $q_2 = 15$. Comparing with the previous section, we find that firm 1 (the leader) earns more profit than it did under the simultaneous Cournot game, but firm 2 (the follower) earns less profit than in the Cournot equilibrium.

The Stackelberg equilibrium is a subgame perfect equilibrium. Are there other Nash equilibria that are not subgame perfect equilibria? Sure – Suppose that firm 2's strategy is to produce the monopoly output $q_2 = 30$ as long as $q_1 = 0$, but threatens to produce $q_2 = 1,000,000$ (driving the market price negative) if firm 1 produces any $q_1 > 0$. If firm 1 actually believed this threat, then it would produce $q_1 = 0$ since it would lose a lot of money otherwise. However, this is not a credible threat by firm 2 since it would lose lots of money as well by producing $q_2 = 1,000,000$. Thus, this constitutes a Nash Equilibrium, but not a subgame perfect equilibrium since it relies on firm 2 making a non-credible threat off of the equilibrium path.