Forward Induction

Adam Brandenburger*

Version 01/04/07

1 Introduction

The term *backward induction* obviously reflects the fact that the backward-induction algorithm starts at the end of the tree and then works backwards. In a sense, past behavior is inferred from future behavior. This suggests that there might also be an opposite phenomenon, in which future behavior can be inferred from past behavior in some way. This is so. The concept is called *forward induction* (not surprisingly!). A player looks back up the tree, to take account of moves that another player could have made but didn't make, to try to predict what that player will do subsequently. That is, the player tries to 'induce forwards' from other players' past behavior to their future behavior.

In this note, we'll look at forward induction. We'll start with some examples, which we'll analyze heuristically. Then we'll give a formal solution concept.

2 Examples

Example 1 [1] *There are two players, labelled A and B. Player B is an incumbent in a certain market, while player A is a potential entrant. If player A enters, then each player must choose between adopting a peaceful strategy (Peace) and a warlike strategy (War). The game is depicted in Figure 1 below.*

Player A can enter the market in which B operates only if A stops serving its current market. If A decides instead to stay in its current market, it will earn a payoff of 1*. Player B will then earn a payoff of* 5*. The full game tree is in Figure 2.*

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

[1] From "Signalling Games," by Adam Brandenburger and Harborne Stuart, teaching material, 02/99.

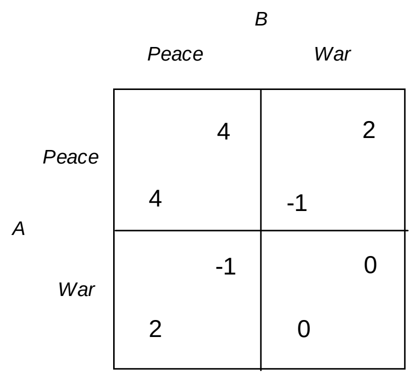

Figure 1

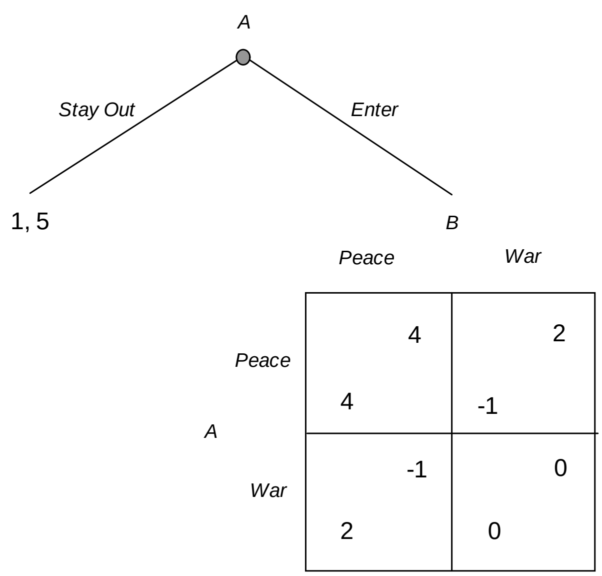

Figure 2

(In Figure 2, player A must first choose between Stay Out and Enter. If player A chooses Stay Out, the game is over. If player A chooses Enter, player B observes this. Each player must then choose between Peace and War, without knowing which choice the other is making.)

How do you think this game might be played?

Example 2 *There are two players: an author and a publisher. The author has to choose between putting a lot of effort into the writing of her book, and making less effort. The publisher has to choose between aggressively promoting the book, and leaving the book to sell itself. Each player has to decide without knowing the other player's decision. Some plausible payoffs are given in the matrix in Figure 3 below.*

Note that this is a coordination game. The best outcome for both players is if the author works hard and the publisher promotes aggressively. But, if either player makes the opposite choice, then so should the other; and both end up worse off.

		Publisher	
		Promote	*No*
Author	*Work*	3 (author), 3 (publisher)	-1 (author), 0 (publisher)
	No	0 (author), -1 (publisher)	0 (author), 0 (publisher)

Figure 3

Now add a prior choice available to the publisher, which is to pay a so-called advance to the author. The advance, if paid, results in a payoff of 1 *to the author (over and above her existing payoffs), and costs the publisher the same amount. The new game tree is in Figure 4.*

What inference do you think the author might draw from a decision by the publisher to pay an advance? What response–work hard or not–do you think the author would then make? What are the implications for the publisher? Is there a sense in which the publisher, in deciding to pay an advance, is sending a signal to the author? Is sending this signal a good idea?

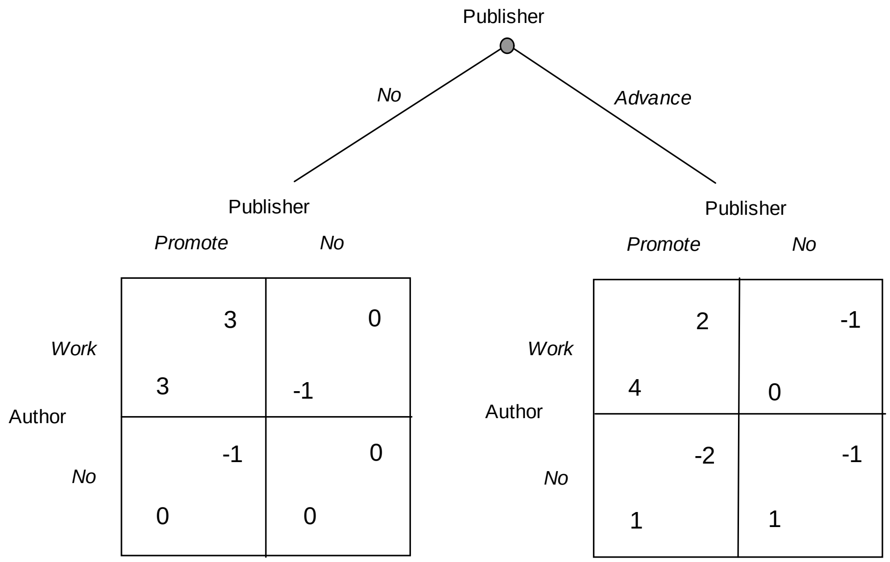

Figure 4

Example 3 *This is a signalling game of the kind pioneered by A. Michael Spence in his book Market Signalling (Harvard Univ. Press, 1974). The specific game is similar to that in "Limit Pricing and Entry under Incomplete Information," by Paul Milgrom and John Roberts, Econometrica, 50, 1982, 443-459.*

We begin with an informal description of a game between an incumbent seller and a potential entrant. Thus, there is an incumbent player, labeled A, selling to a certain buyer at a cost of $40. *There is also a second player, labeled B, that is considering entering the business to compete for the buyer.*

The game unfolds in three stages. In the first stage, player A sells to the buyer–at a price that player B observes. In the second stage of the game, player B decides whether or not to enter the business. If it enters, it incurs an irrecoverable entry cost of $1, *and then has a* $50 *cost of serving the buyer.*

If player B decides to enter, then, in the third stage, players A and B compete for the buyer's next purchase. If, in the second stage, player B decides not to enter the business, then, in the third stage, player A sells again to the buyer without facing any competition.

Before entering the business, player B does not know player A's cost. In fact, prior to the first stage of the game, player B believes that A's cost might be $40 *or* $60, *and assigns a probability of* 1/2 *to each possibility.*

The highest price that the buyer can be charged is $100. *But, if players A and B compete for the buyer, then the price is driven below this–in fact, to a shade less than A's cost or B's cost, whichever is the higher. Thus, if player A's* *cost is* \$40 *(as is actually the case), the resulting price will be a shade less than* \$50*, and A will secure the buyer's business. If A's cost is* \$60 *(which player B believes possible), then the resulting price will be a shade less than* \$60*, and B will secure the buyer's business.*

In the first stage of the game, player A must decide whether to charge the monopoly price of \$100*, or a 'deterrent' price of* \$55*.*

A game tree depicting this situation is drawn in Figure 5.

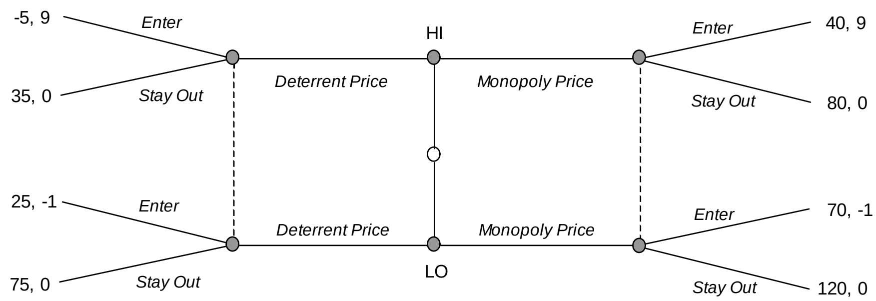

Figure 5

Here, the open node denotes an initial 'chance' move that, with 1/2 *probability, assigns a cost of* \$40 *(LO) to player A, and, with* 1/2 *probability, assigns a cost of* \$60 *(HI) to player A.*[2] *Knowing its cost, player A then decides whether to charge the monopoly price or the deterrent price. The dotted lines are information sets indicating that player B gets to observe the price A charges, but not the level of A's cost. Thus, if player A charges the monopoly price, player B knows only that it is at the upper-right-hand solid node or the lower-right-hand solid node (but not which of the two). Player B must then decide whether to enter or stay out. Likewise, if player A charges the deterrent price, player B knows only that it is at the upper-left-hand solid node or the lower-left-hand solid node (but not which of the two). In this case, too, player B must decide whether to enter or stay out. Finally, the left-hand number at each endpoint is the resulting payoff to player A, the right-hand number is the resulting payoff to player B.*

To analyze the game, begin by verifying that the payoffs in the game tree accord with the verbal description of the game given above. Now, ask: If player A charges the 'deterrent' price, will player B, in fact, be deterred (i.e. decide to stay out)? Is charging the 'deterrent' price a good idea for player A?

[2] The initial chance move (by Nature, if you like) is a new feature, beyond our definition of a tree in the note "Game Trees." The present game is an example of a class of games in which Nature first selects a so-called *payoff type* for each player. (Here, the types are LO or HI for player A; there is no choice of type for player B.) Each player is then informed of its own payoff type, but not the payoff types of the other players. Then, the players decide on their moves. These games are often referred to as *Bayesian games*.

3 (Iterated) Weak Dominance

We now introduce a new solution concept, which will turn out to be very useful in a formal analysis of the preceding examples.

We follow the set-up and notation in the note "Dominance and Iterated Dominance." In particular, then, for any finite set X, we let $\Delta(X)$ denote the set of probability measures (distributions) over X. Also, we now write $\Delta^0(X)$ for the subset of $\Delta(X)$ consisting of those measures that give strictly positive probability to every point $x \in X$.

We start by strengthening our earlier definition of rationality (Definition 4 in "Dominance and Iterated Dominance"). The new definition says that Ann considers each of Bob's strategy choices possible; she rules nothing out.

Definition 1 *A strategy $s^a \in S^a$ is **rational** if there is a $\sigma^b \in \Delta^0(S^b)$ such that $v^a(s^a, \sigma^b) \geq v^a(r^a, \sigma^b)$ for all $r^a \in S^a$.*

(Note that, as before, we state the definition in terms of one of Bob's mixed strategies σ^b. But also as before, in the interpretation we are going to think of this mixed strategy as really being a probability distribution held by Ann about Bob's (definite) choice of strategy.)

Next, we relate this stronger definition of rationality to a concept of dominance that is likewise stronger than our earlier one (Definition 3 in "Dominance and Iterated Dominance"). The stronger concept is (rightly) called *weak dominance*, as opposed to our earlier *strong dominance*. Do not be confused!

Definition 2 *A strategy $s^a \in S^a$ is **weakly dominated** (or **inadmissible**) if there is a $\sigma^a \in \Delta(S^a)$ such that*

$$\begin{aligned} v^a(\sigma^a, s^b) &\geq v^a(s^a, s^b) \text{ for every } s^b \in S^b, \\ v^a(\sigma^a, s^b) &> v^a(s^a, s^b) \text{ for some } s^b \in S^b. \end{aligned}$$

Proposition 1 *A strategy $s^a \in S^a$ is rational if and only if it is not weakly dominated (admissible).*

The proof is very similar to the proof of the analogous Proposition 1 in "Dominance and Iterated Dominance." (Try the proof as an exercise.)

Continuing to parallel what we did before, we now define an elimination procedure called *iterated weak dominance* or *iterated admissibility*. Verbally, the procedure is:

Step 1: For each player, delete all strategies of that player that are inadmissible in the game. This defines a 'reduced' game.

Step 2: For each player, delete all strategies of that player that are inadmissible in the reduced game. This defines a further reduced game.

...

The procedure continues until no further deletion is possible, and a 'residual' game is produced. The strategies in this residual game are called the *iteratively admissible* strategies.

Here is the formal definition. Let $S^a_0 = S^a$ and $S^b_0 = S^b$, and define S^a_m and S^b_m inductively by

$$S^a_{m+1} = \{s^a \in S^a_m : s^a \text{ is admissible in the game } \langle S^a_m, S^b_m, \pi^a|(S^a_m \times S^b_m), \pi^b|(S^b_m \times S^a_m)\rangle\},$$

$$S^b_{m+1} = \{s^b \in S^b_m : s^b \text{ is admissible in the game } \langle S^a_m, S^b_m, \pi^a|(S^a_m \times S^b_m), \pi^b|(S^b_m \times S^a_m)\rangle\},$$

Note that since S^a and S^b are finite, there is an M such that $S^a_m = S^a_M$ and $S^b_m = S^b_M$ for all $m \geq M$. (Check why this is true!)

Definition 3 *A strategy $s^a \in S^a_M$ or $s^b \in S^b_M$ is called **iteratively admissible**.*

4 Examples Revisited

The iteratively admissible strategies can be thought of as the strategies that can be played when both players are rational–in the strengthened sense of this note–and there is common belief of this notion of rationality. (For an exact treatment of this idea, see "Admissibility in Games," by Adam Brandenburger, Amanda Friedenberg, and H. Jerome Keisler, at www.stern.nyu.edu/∼abranden.)

Happily, iterated admissibility gives exactly the desired outcome in the examples of Section 2 in this note. (Verify this!) Iterated admissibility also gives the backward-induction outcome in perfect-information trees. It thus gives us a unified solution concept for all game trees.