

An Introduction to Game Theory Using the GamePlan Software[†]

By

Jean-Pierre P. Langlois

Email: Langlois@math.sfsu.edu

Web: <http://userwww.sfsu.edu/~langlois/>

November 9, 2005

Abstract

This short paper introduces the extensive and normal forms of a game and discusses Nash, subgame perfect, and perfect Bayesian equilibria. It illustrates these concepts with typical games such as Chicken, Selten's Horse, the Dollar Auction, and the Prisoner's Dilemma in both its one-shot and its repeated versions.

[†] Copyright © 2005, Jean-Pierre P. Langlois. All rights reserved.

The Concepts of Game Theory

The Extensive Form

Formally, a game is a structure made up of five basic types of objects: players, states (also called nodes), moves (also called choices), information sets, and payoffs (also called utilities). Some additional concepts such as chance moves and discounting appear in more advanced game structures.

There needs to be at least two players in a game, at least one turn of play represented by a node or an information set, and at least two choices. Some moves end the game and are called “final” while others yield to a next turn. Final moves must yield “outcomes” in the form of payoffs to all the players involved in the game. Payoffs are real numbers. The more a player prefers an outcome, the higher the payoff to that player.

Figure 1 was produced with the *GamePlan* software and shows what is arguably the simplest possible game in its so-called “extensive form.”

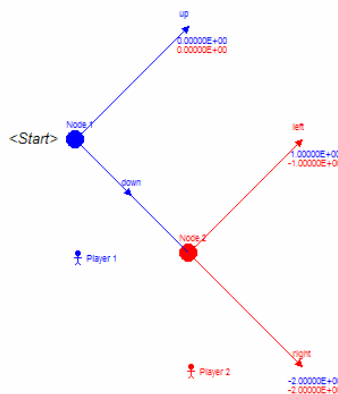


Figure 1: The Simplest Game

Simple games such as this one are organized in a “tree form” with a well defined root (here Node 1) where the game is assumed to start. *GamePlan* uses color coding to recognize easily how objects are associated. Player 1 here is coded in blue as illustrated by the little blue icon outlining a human form. Everything colored in blue in the picture thus belongs to Player 1. Since Node 1 is blue, Player 1 owns it and therefore has the first turn. Player 2 is in red and owns Node 2. The arrows represent the moves available to the player who has the turn. At Node 1, Player 1 can choose “up” or “down” since these are blue and clearly originate from Node 1. “up” ends with a pair of real numbers written in standard scientific notation. Again, the color says who gets what if that move is chosen. 0.00000E+00 (in blue) is a complicated way of writing that the payoff to Player 1 is 0 (and similarly to Player 2). The choice of “down” by Player 1 ends in Node 2 at which Player 2 has the turn since the node is red. Player 2’s available moves are “left” and “right” and are both final, each ending in a pair of payoffs. The game of Figure 1 has no

information set, meaning that the players know exactly where they are in the game tree when it is their turn to play.

Figure 2 was also produced with *GamePlan* and shows the simplest possible structure with an information set. With the given payoffs, this game is known as Chicken.

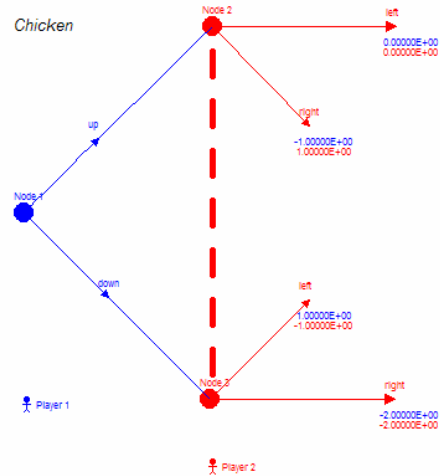


Figure 2: A Game with an Information Set

Now, Player 2 has two nodes instead of one and they are joined by a thick red broken line that denotes an information set in *GamePlan*. This means that Player 2 doesn't know whether he is at Node 2 or Node 3 when his turn comes. Put another way, Player 2 cannot observe Player 1's choice at Node 1, whereas he could infer it in Figure 1. This also means that the exact same choices (left or right) must be available at both nodes of the information set. Otherwise, Player 2 would know where he is according to what choices are available. For all strategic purposes, it is as if both sides make their choices simultaneously in Figure 2. So, the presence or absence of an information set can translate into a timing constraint.

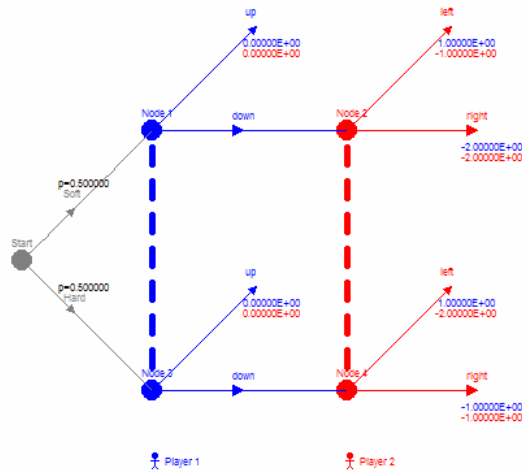


Figure 3

Information sets do not just indicate timing issues. Figure 3 shows what is usually called an “incomplete information” game.

First, there is an additional gray node labeled “Start” and two gray moves (called soft and hard) with “ $p=0.500000$ ” written above their name. Gray nodes indicate Chance (or Nature) turns and the moves issuing from them must have some fixed probability distribution. Nature is *not* a player in the game since she has no objectives (payoffs) and does not decide the probability she uses to do this or that. But she influences decisively the unfolding of the game. Before any of the actual players have a turn, Chance flips a fair coin and decides whether the game will follow the upper or the lower branch. The blue and red information sets then mean that neither side knows which of the two branches they are in. But the information sets do not indicate simultaneous play here. Indeed, Player 2 only has a say if Player 1 doesn’t choose “up” at her turn. So he can infer what Player 1 did *before* making his choice.

Games must have at least two real players (Nature doesn’t count) but often have more. *GamePlan* allows up to three players. Figure 4 shows a 3-player game.

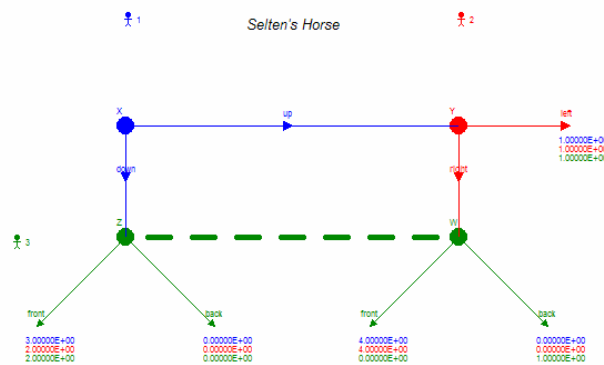


Figure 4: Selten’s Horse

This is known as Selten’s horse (from the name of Nobel Laureate Reinhart Selten). The interesting feature here is that 3 doesn’t know whether 1 or 2 played when he does, and he may never play at all if 1 chooses up and 2 chooses left, although he will know what happened when handed his payoff of 1.0.

In general, each player has several turns in a game depending on prior developments. In fact, the number of turns may even be subject to the players’ decisions or to Chance. An interesting “real life” example that is a favorite of game theorists is the so-called Dollar Auction. There are variations on the theme but I will describe a simple version that can be illustrated in *GamePlan*: a professor starts her class by announcing that she will auction a \$10 bill to the highest bidder. The last bidder gets the \$10 bill for whatever he or she bid last. However, there is a twist: the next-to-last bidder also pays whatever he or she bid last but gets nothing.

In general the bids can be any multiple of a dollar but to simplify let us assume that they can only be exactly one dollar at a time. To simplify even further, let us assume that the professor picks (only) two students who will be allowed to bid, with Student 1 going first. Student 1 can either stop, getting nothing, in which case we assume that the professor keeps the \$10 bill. Or he can bid \$1 and let Student 2 decide what to do next. If Student 2 chooses to stop, Student 1 gets \$10 minus the \$1 he bid for a net of \$9. If Student 2 bids another \$1, it is Student 1's turn to choose again. From then on, each time a student bids he adds \$1 to what the professor gets for her \$10 bill and adds \$2 to his own previous bid. The game is pictured in Figure 5.

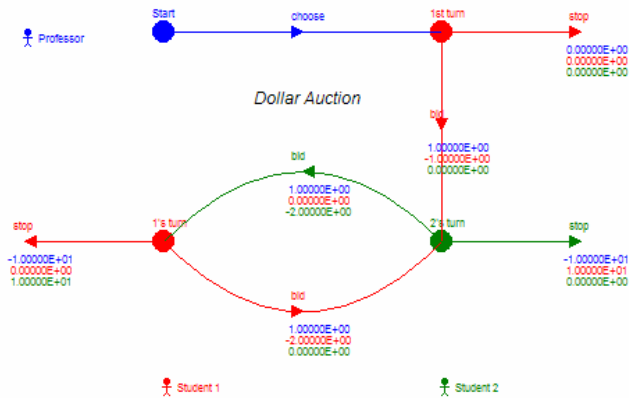


Figure 5: A Simplified Dollar Auction

This game structure illustrates two important facts: first, a game is not necessarily played on a tree since Figure 5 involves a loop that can be visited any number of times. Second, non-final moves (the bid moves) can involve payoffs to the players. Such “transient” payoffs simply accumulate according to developments until the game ends.

The Normal Form

All four above examples use the extensive form of a game. But game theory was initially popularized in the social sciences using the so-called “normal” (also called “strategic”) form. It begins with the concept of “strategy” which is simply a gameplan for a player. In the games where each player has a single node or information set, as in the first four above examples, it is easy to confuse the concept of strategy with that of move. In Figure 1, Player 1 can choose only up or down and these are her possible *pure* strategies as well. In Figure 2, Player 2 also has only two pure strategies (left or right) although these involve four possible moves depending on what Player 1 does. In Figure 3 the situation is similar with again two pure strategies per player although each player technically owns four moves. Figure 4 is hardly different from that point of view. But Figure 5 can be misleading: in reality there are infinitely many potential moves for each player. So, a player could for instance plan to bid for up to five times and then stop if he ever gets there. But he could also decide to bid up to one hundred times. The number of pure strategies for the game of Figure 5 is therefore infinite.

Traditionally, the normal form is presented as a table. Table 1 is the normal form of the game of Figure 1.

		Player 2	
		left	right
Player 1	up	0.00000E+00 0.00000E+00	0.00000E+00 0.00000E+00
	down	1.00000E+00 -1.00000E+00	-2.00000E+00 -2.00000E+00

Table 1: The Normal Form of Figure 1

The *strategy* “up” (by opposition to the *move* “up”) now yields the payoff 0 to each side regardless of what Player 2 chooses. The upper boxes in Table 1 therefore list the same pair of payoffs of zero. And since the strategy “down” makes Player 2’s choice relevant, the lower two boxes show different payoffs.

The normal form corresponding to Figure 2 is shown in Table 2.

		Player 2	
		left	right
Player 1	up	0.00000E+00 0.00000E+00	-1.00000E+00 1.00000E+00
	down	1.00000E+00 -1.00000E+00	-2.00000E+00 -2.00000E+00

Table 2: The Normal Form of Figure 2

The normal form corresponding to Figure 3 is represented in Table 3.

		Player 2	
		left	right
Player 1	up	0.00000E+00 0.00000E+00	0.00000E+00 0.00000E+00
	down	1.00000E+00 -1.50000E+00	-1.50000E+00 -1.50000E+00

Table 3: The Normal Form of Figure 3.

It requires a bit more thinking. Now the players' choices can yield different results depending on Nature's move. "up" gives each side 0 in either case and that is the payoff of "up" in the upper boxes of Table 3. Also, "down" and "left" yield 1 to player 1 in either case as reflected in the lower left box. But it yields -1 in the soft case and -2 in the hard case to Player 2, each with probability 0.5. The resulting expected red payoff -1.5 is therefore entered in the lower left box. The lower right box follows the same logic.

The normal form of the game of Figure 4 cannot be represented by a two-dimensional table. Instead, it is the three dimensional table illustrated in Table 4.

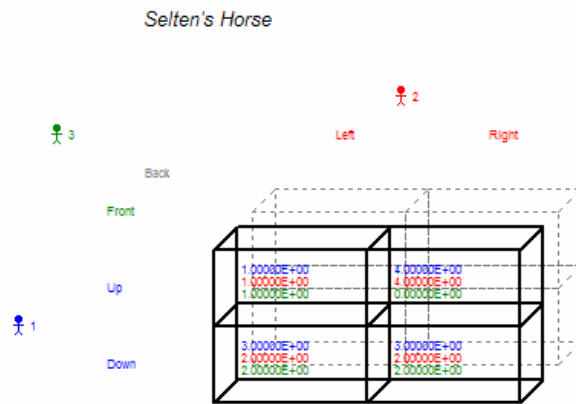


Table 4: The Normal Form of Selten's Horse

It shows only the payoffs when 3 uses "front". A similar table (the broken lines) must be filled in when 3 uses "back".

The normal form of Figure 5 would of course involve infinitely many rows and columns and therefore cannot be represented.

Solution Concepts

The origins of formal game theory can be traced back to Augustin Cournot in the 19th century. A more recent origin is John von Neumann in a 1928 seminal article on zero-sum games (what one side wins is what the other loses). But the main figure is undoubtedly John Nash who, in a 1950 seminal article introduced the modern concept of (Nash) equilibrium: a choice of one strategy per player such that no one can benefit by deviating unilaterally (in an expected payoff sense). Most importantly, Nash proved the existence of such equilibria in all games, provided the randomization of strategies is allowed. Usually, randomization of strategy is equivalent to the randomization of the corresponding moves.¹ An equilibrium without any randomization is called "pure" and one that involves probabilities is called "mixed".

¹ This is true when the extensive form has the property of "perfect recall". Loosely speaking, this means that a player's information sets are such that he never forgets information previously held.

Although existence is ensured, uniqueness is far from the norm. The game of Figure 1 has in fact numerous Nash equilibria. They can be described in both extensive and normal forms. For instance

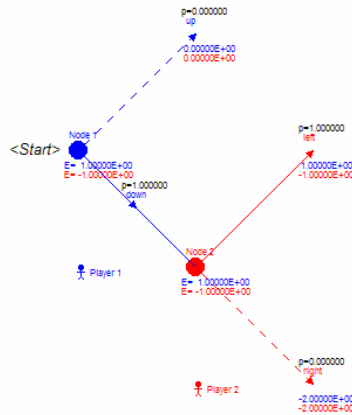


Figure 6: A Nash Equilibrium

In *GamePlan* an equilibrium is shown by giving the probability of each move in the form “p=” above the name of the move. Moves with probability zero are pictured with a dotted rather than solid line. In Figure 6, both “down” and “left” have probability p=1. Each node also has a pair of numbers beginning with “E=” (for expected payoff). At Node 2, for instance, “E=1.00000E+00” in blue is a complicated way of saying that Player 1 expects a final payoff of 1 by choosing “down” in the expectation that Player 2 will choose “left”. Of course, this is better for Player 1 than choosing “up” which yields the actual payoff of 0. In turn, Player 2 finds it best to choose “left” for a red payoff of -1 rather than “right” that yields the red -2. Indeed, the expected payoffs for the entire equilibrium are shown with “E=” under the start node. Clearly, neither side would benefit from any deviation from this gameplan.

The same equilibrium can be viewed in the normal form of Table 5. Here again the probability p=1 is attached to the strategies “down” and “left” and the other two (“up” and “right”) are grayed to indicate zero probability. The expected payoffs are clearly 1(for blue) and -1 (for red) in the bottom-left box.

		Player 2	
		left p= 1.000000	right p= 0.000000
Player 1	up p= 0.000000	0.00000E+00 0.00000E+00	0.00000E+00 0.00000E+00
	down p= 1.000000	1.00000E+00 -1.00000E+00	-2.00000E+00 -2.00000E+00

Table 5: A Nash Equilibrium

But *GamePlan* actually shows two other equilibria in the normal form. Here is one:

		Player 2	
		left p= 0.000000	right p= 1.000000
Player 1	up p= 1.000000	0.00000E+00 0.00000E+00	0.00000E+00 0.00000E+00
	down p= 0.000000	1.00000E+00 -1.00000E+00	-2.00000E+00 -2.00000E+00

Table 6: Another Nash Equilibrium

Here, “p=1” is attached to “up” and “right”. Indeed, Player 1 would lose by unilaterally switching to “down” since that would yield the blue payoff -2 instead of 0 in the bottom-right box. Player 2 would not lose but would not benefit either by switching unilaterally from “right” to “left” since this would yield the same red payoff of zero in either top boxes. The Nash equilibrium concept only requires that there be no benefit from switching. It doesn’t require that there be an actual loss for either side.

The equilibrium of Table 6 can also be seen in the extensive form in Figure 7. Here again “up” and “right” have probability “p=1” and the expected payoffs under the start node are both “E=0”. But there is something a little odd here: at Node 2 Player 2 is planning to choose “right” although, if he ever found himself there, he should plan on choosing “up” which yields a strictly greater red payoff of -1 instead of -2. The fact is that he is not *planning* to ever reach Node 2 in this Nash equilibrium. So, whatever he plans for Node 2 has no effect on his own expected payoff at Node 1. But it does have an effect on Player 1’s expected payoff: indeed, by planning “right” at Node 2 he creates the blue expectation “E=-2” at Node 2 and therefore justifies the superior blue choice of “up” at Node 1. It is a bit as if Player 2 is deterring Player 1 from choosing “down” by an implicit threat of choosing “right”.

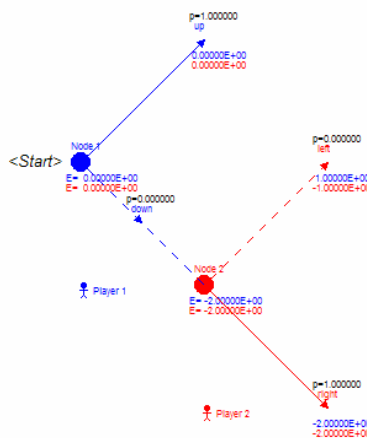


Figure 7: The Nash Equilibrium of Figure 6

But that threat is flimsy! There is no way Player 2 would *rationally* choose “right” if Player 1 decided to call the bluff by choosing “down”. The Nash equilibrium can make perfect sense in the normal form that *erases* completely all timing issues but it is not always convincing in the extensive form. This annoying feature led to refining the concept of Nash into that of “perfect” equilibrium (in the later work of Selten and Harsanyi). These ideas will be discussed in the third section of these notes. The equilibrium of Figure 6 is (subgame) perfect while that of Figure 7 is not. In fact there is a continuum of Nash but non-perfect equilibria in the above game where Player 1 chooses “up” and player 2 chooses “right” with probability at least 0.5.

The game of Figure 2 (and Table 2) is known as “Chicken” and has exactly three Nash equilibria. Two are pure and as follows: “up and right” yielding the payoffs -1 for blue and 1 for red, and “down and left” yielding 1 for blue and -1 for red. There is also a mixed equilibrium with each side using a 0.5 probability on each strategy.

The game of Figure 3 (and Table 3) also has several Nash equilibria: “down and left” and a continuum from “up and right” to “up together with left or right with 50% chances”.

The situation is similar in Figure 4 (and Table 4) with Nash equilibria such as “up, left and back”, “down, left and front” and “down, front, and 50% chances of right or left” as well as many other combinations. The second one is illustrated in Figure 8. It suffers from the same credibility problems identified in Figure 7: the choice of “left” by 2 would not be rational at Node Y if play ever reached it. But the flimsy threat of “left” makes 1 choose “down” in Nash equilibrium.

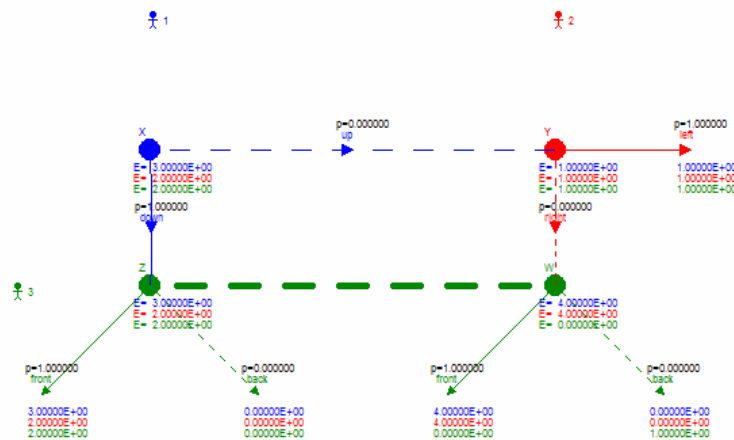


Figure 8: A Non-Perfect Nash Equilibrium

Finally, *GamePlan* finds three equilibria in the game of Figure 5, two of which are pure and one that involves probabilistic moves. This last one is shown in Figure 9.

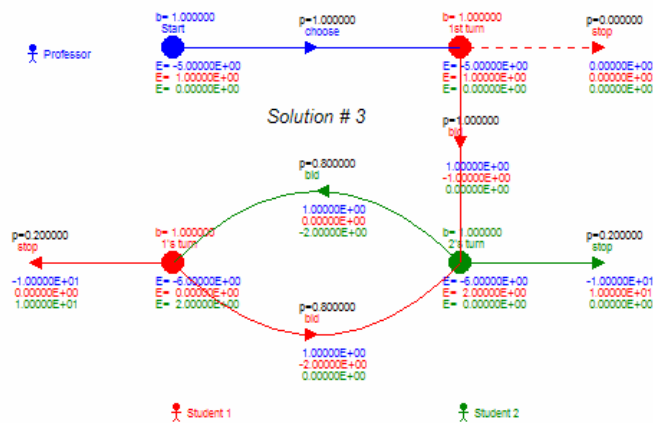


Figure 9: An Equilibrium of the Dollar Auction

In this case, Student 1 bids with certainty at his first turn and both sides continue bidding up with 80% chance turn after turn. The rationale for this behavior is easy to explain by looking at the picture: note that red's expected payoff at the green "2's turn" node is $E=2$. But by bidding from node "1's turn" red pays \$2 as reflected in the -2 red payoff under that move. So, red is expecting $-2+2=0$ from making that bid. But that is exactly his actual payoff from stop. So, red is indifferent between "bid" and "stop" and can choose either with the given probabilities. In turn, the red expected $E=2$ at node "2's turn" is easily explained: green chooses "stop" that yields the red payoff 10 with probability 0.2 and "bid" that yields red an expected $0+0=0$ with probability 0.8 for an expected total of $E=2$. The situation is entirely symmetric for green. Perhaps this equilibrium explains the empirical observation that students engaged in the game often bid up beyond the very value (\$10) of the prize. It also illustrates one important fact: equilibria are always *forward* looking. As the game unfolds in Figure 9 the two sides accumulate losses as they bid repeatedly. But their rational behavior remains driven by their expectation of the future. The past costs are *sunk* and do not influence their gameplan for the future.

Homework

1. The game of Figure 2 was interpreted as one of simultaneous play. So, the two players' turns should be interchangeable without consequences. Construct an extensive form equivalent to Figure 2 and Table 2 where Player 2 has the start node.
2. Modify the game of Figure 3 so that there is no blue information set. Then write its normal form (hint: Player 1 now has four strategies).
3. Write the normal form of the game of Figure 3 with the probabilities of "soft" and "hard" changed to 0.4 and 0.6 respectively.
4. Replace the \$10 bill by a \$20 bill in the game of Figure 5. Verify (without using *GamePlan*) that the repeated bids with probability $p=0.9$ form an equilibrium.

Advanced Game Structures and Solution Concepts

It is significant that the first Nobel Prize awarded (in 1994) to game theorists went to John Nash, John Harsanyi, and Reinhard Selten. Nash developed the very concept of

game equilibrium while Harsanyi and Selten refined it to remedy some of its weaknesses: Selten developed the idea of “subgame perfect” equilibrium and Harsanyi that of “Bayesian” equilibrium, two modern standards in game theory.²

Subgame Perfect Equilibrium (SPE)

Some limitations of the Nash equilibrium concept were already outlined in Figure 7: the planned move “right” at Node 2 is not credible because it would not be optimal if Player 1 chose “down”. Selten observed that the structure beginning at Node 2 with the two moves “left” and “right” actually forms an extremely simple game all by itself: although Player 1 doesn’t have a move there she does receive a payoff. That structure is called a “subgame” of the entire game beginning at Node 1. And this subgame has a trivial Nash equilibrium where Player 2 chooses “left”. So, Selten proposes to consider the set of all possible subgames of a game (including the whole game itself).

Formally, if a Nash equilibrium of the whole game translates into a Nash equilibrium in every subgame it is called “subgame perfect”.

The equilibrium of Figure 7 fails that test but that of Figure 6 passes it with flying colors.

Clearly, the SPE concept is only applicable to extensive form games.³ But does it always resolve the credibility problems of the Nash equilibrium? Selten himself gave a counter-example that answers the question negatively. It is the game of Figure 4. The trouble is that it admits no subgame other than itself: indeed, one might be tempted to say that the “game” beginning at Node Y is a subgame. But the move from Y to W yields to part of the information set {Z,W} which cannot be broken without changing the nature of the game. In other words, a “subgame” must be a self-contained game all by itself. One cannot cut out any part that doesn’t strictly precede it. So, the inconsistency pointed out in Figure 8 is not resolved by the SPE concept.

The game of Figure 3 offers a similar hurdle: the only subgame is the entire game itself so that any Nash equilibrium is trivially a SPE.

But how does this SPE concept apply to Figure 5? The figure is in fact misleading from the subgame viewpoint: a more adequate, although incomplete, picture would be Figure 10. The structure repeats indefinitely, as suggested by the red dotted line at the bottom. But one can see that there is a subgame starting at any node of the entire tree. So, the SPE concept applies and should resolve any credibility issue. In fact, the equilibrium of Figure 9 translates without any change to Figure 10. It only has the unusual property that equilibrium play at any node (but the first two) is always to bid with 80% probability. In other words, the prior developments don’t matter and one can just as well look at the graph of Figure 9 and decide that only two states of the game need be considered: whether it is red’s or green’s turn to play.

² The second Nobel Prize given (in 2005) to game theorists went to Robert Aumann and Thomas Schelling. Aumann has been a main figure in the developments of “repeated games” while Schelling has applied game theoretic concepts and methods to the analysis of conflict and deterrence.

³ Selten also developed a “trembling hand” equilibrium concept for the normal form.

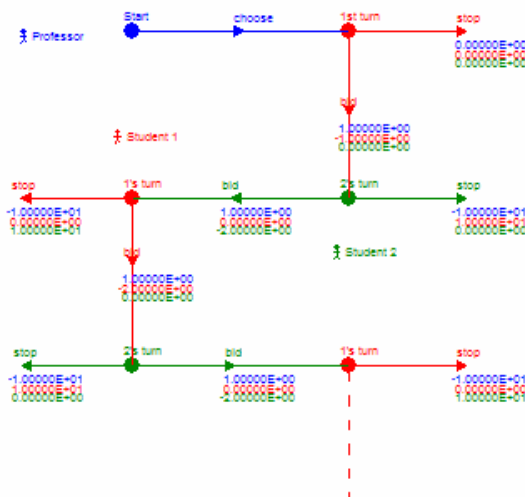


Figure 10: A Tree Form of Figure 3

An equilibrium where all kinds of prior developments are subsumed into a few states is called “Markov perfect”. A Markov perfect equilibrium (MPE) in such a repeating game is always an SPE with a simple structure.

Perfect Bayesian Equilibrium (PBE)

Incomplete information games such as that of Figure 3 are mostly ill suited to the SPE concept. In essence, they arise when the players do not know whether they are playing one game or another, something that is decided by chance at the start. So, any “cut” would presumably cut an information set, something not allowed.

Harsanyi is at the origin of the Bayesian approach to this kind of problems. It is well illustrated with the following variation on the theme of Figure 3.

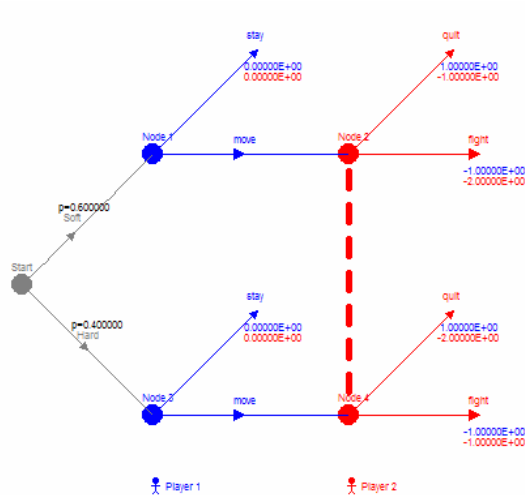


Figure 11: One-Sided Incomplete Information

Here, Player 1 *knows* what game she is playing but Player 2 doesn't, so that information is called "one-sided." The standard interpretation is that Player 1 has two "types" (here labeled "soft" and "hard") and that the outcome of the two sides' choices will depend on what type Player 2 was actually playing against. As always in such cases, the initial probabilities (60% soft, 40% hard) called "prior beliefs" are known to both sides.

Again, there is no (proper) subgame but the whole game itself. The normal form will involve four strategies for Player 1 and two for Player 2 and yield numerous Nash equilibria.

Harsanyi's insight is that Player 2 would not only act according to the initial chances that he would be playing in the upper or the lower part of the tree. He would also interpret Player 1's move as indicative of her type and could modify his *beliefs* about who he is really facing accordingly. The concept of beliefs is formalized as a probability distribution over the nodes of each information set.⁴ At his turn, Player 2 should decide rationally whether to choose "left" or "right" and this depends on where (Node 2 or 4) he believes to be. If, for instance, he believes to be at Node 2 with 60% chances as the prior beliefs would indicate he should choose "quit" with certainty, a choice with an expected payoff of -1.4 rather than -1.6 for "fight". But if Player 2 chooses "quit" with certainty Player 1 is strictly better off choosing "move" that yields an expected payoff of 1 rather than "stay" that only yields 0 for either type. Figure 12 shows this equilibrium.

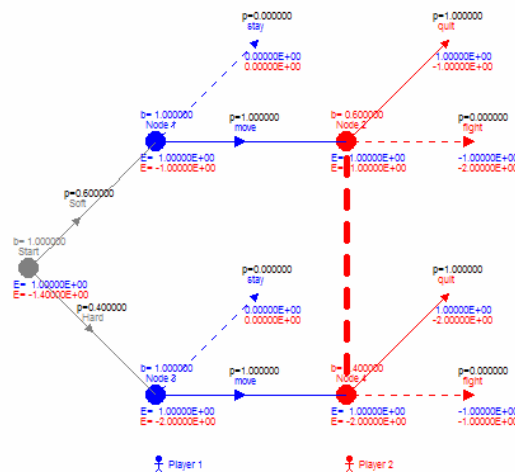


Figure 12: An Equilibrium

The notation "b=0.6" above Node 2 shows Player 2's belief (also known to Player 1) that he is there at his turn of play. But could he entertain different beliefs and therefore make different choices? Figure 13 shows precisely such a case.

⁴ A node that is not part of any multi-node information set is also called a "singleton" information set and must always have belief 1.

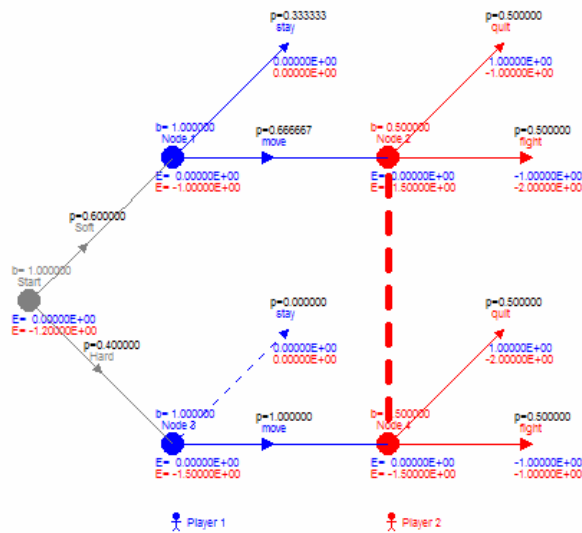


Figure 13: Another Equilibrium

Now, Player 2 believes to be at Node 2 with only 50% chances. As a result he is indifferent between “quit” and “fight” and can play them with 50% probabilities. The effect is that the *soft* Player 1 becomes indifferent between “stay” and “move” at Node 1 and can play them with 1/3-2/3 probabilities. But how did the beliefs change from the prior 60% on soft to the “posterior” 50%? This is precisely Harsanyi’s insight: Player 2 applies Bayes Law of conditional probabilities as follows:

$$P(\text{soft}|2\text{'s turn})P(2\text{'s turn})=P(2\text{'s turn}|\text{soft})P(\text{soft})=P(\text{move}|\text{soft})P(\text{soft})$$
 Observing that $P(2\text{'s turn})=P(\text{move}|\text{soft})P(\text{soft})+P(\text{move}|\text{hard})P(\text{hard})$, this yields the standard Bayesian updating formula

$$b \text{ (at Node 2)} = \frac{P(\text{soft}|2\text{'s turn})}{P(\text{move}|\text{soft})P(\text{soft}) + P(\text{move}|\text{hard})P(\text{hard})}$$

Since $P(\text{move}|\text{soft})=2/3$, while $P(\text{move}|\text{hard})=1$, according to Figure 13, one gets

$$b = \frac{(2/3)(0.6)}{((2/3)(0.6) + (1)(0.4))} = 0.5$$

GamePlan does these calculations automatically but it is important to understand the underlying logic. Both figures 12 and 13 show perfect Bayesian equilibria (PBE) that are clearly far from unique. Indeed, *GamePlan* provides three other PBEs, one of which is pictured in Figure 14.

Note that Player 2 believes to be at Node 4 with certainty (if his turn comes) and chooses “fight” accordingly. This deters both types of Player 1 to move and prevents Player 2 from making any actual decision. Bayesian updating here plays no role whatsoever because the red information set is never reached in equilibrium play. In that case, any

distribution of beliefs on the red nodes is allowed in PBE. One says that the red information set is “off the equilibrium path”.⁵

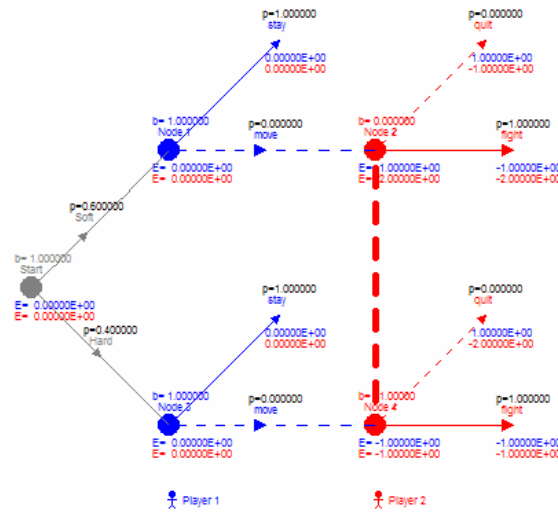


Figure 14: Yet Another PBE

Formally, a PBE is made up of one strategy per type and (common) beliefs at all information sets such that the beliefs are consistent with the strategies by Bayes law (and arbitrary off the equilibrium path) and the strategies prescribe optimal choices at each turn given the beliefs (they are sequentially rational).

The Shadow of the Future

The game of Figure 5, or its infinite tree form in Figure 10, is an example of a more general class of “repeated games”. Such games can either stop when the players decide to or they can stop as a result of a chance move by Nature. Repeated games have been studied extensively and are a main source of advanced game models together with the incomplete information ones.⁶ Aumann gave an extensive survey of theoretical results in 1981.

A major brand of repeated games uses the concept of “discounting” of the future. For instance, in the repeated game of Figure 5, future bid-costs and prospects of eventually winning the auction could be weighed less than present ones. An alternative way of thinking is that Nature may interrupt the game at any time with some probability. For instance, the professor could use a random device to decide whether the auction will go through another turn.

⁵ The arbitrariness of beliefs off the equilibrium path can be criticized as allowing the non-credible threat of “fight” in Figure 14.

⁶ It is possible to work with repeated games of incomplete information but they are usually hard to solve.

Figure 15 shows how discounting works in *GamePlan*: in the upper game fragment “bid” leads to a chance move with probability $p=0.01$ that the game ends with the red player winning the bid. The game therefore continues with probability $p=0.99$. But that means that all developments beyond turn #N are viewed with less likelihood and therefore less impact on rational calculations.

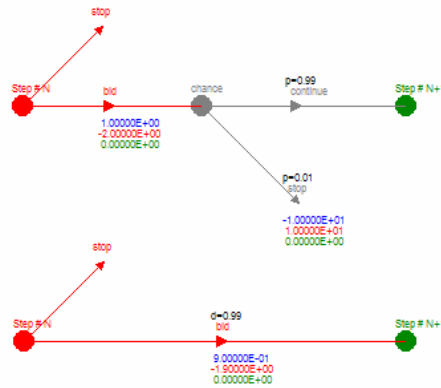


Figure 15: Two Ways of Looking at Discounting

Equivalently, the chance node can be removed altogether and the “bid” move can be assigned the discount factor $d=0.99$. But the transient payoff must be adjusted accordingly: the 0.01 chance of winning \$10 amounts to an expected *instant* win of \$0.1 that must be added to the -\$2 bid for a -\$1.9 payoff (and similarly for the other players). Modifying the game of Figure 5 along these lines has an interesting effect on the mixed solution: the probability of bidding rises! (see homework 4 below).

Repeating a standard “one-shot” game with discounting can have surprising effects on the solution. The most well known example of this fact is the famous Prisoner’s Dilemma, usually introduced in the normal form of Table 7.

	column	
	cooperate	defect
cooperate	0.00000E+00 0.00000E+00	-2.00000E+00 1.00000E+00
defect	1.00000E+00 -2.00000E+00	-1.00000E+00 -1.00000E+00

Table 7: The Classic Prisoner’s Dilemma

As a one-shot game, the solution is trivial: both sides find that “defect” is best regardless of what the other side does. This is called a “dominant” strategy. The single Nash equilibrium “defect-defect” thus yields the payoff -1 to each side. The paradox arises from the fact that both would be strictly better off if they could “trust” each other and

play “cooperate”. Unfortunately, in a simultaneous play situation, this is a recipe for getting exploited by unilateral defection (payoff 1 instead of 0 for the unilateral defector). The Prisoner’s Dilemma has been used as a metaphor for countless situations where individualistic behavior leads to an inferior social outcome.

Is there a way to incite cooperation with the threat of future retaliation? This requires that the game be played more than once. Repeating this game a *finite* number of times makes strictly no difference in the logic of defection: on the last iteration, both sides will defect expecting no future. Thus, they will do so on the next to last turn since defection is a foregone conclusion for the last turn, and so on, back to the very first turn.

However, repeating the game with some *probability* has dramatic consequences. Now future retaliation is always a possibility and perhaps it is ground for cooperation. Writing the extensive form version of Table 1 should by now be elementary (simply edit Figure 2 for payoffs and move names). But writing an extensive form for the *repeated* version is slightly more challenging. It all depends on how many possible developments one wants to distinguish as representative of prior histories. A standard approach is to distinguish four states according to the four possible outcomes of each turn. The resulting extensive form therefore distinguishes four replicas of the one-shot extensive form of Table 7, each corresponding to a possible prior play of the game. Figure 16 shows a particularly interesting SPE in the resulting picture.

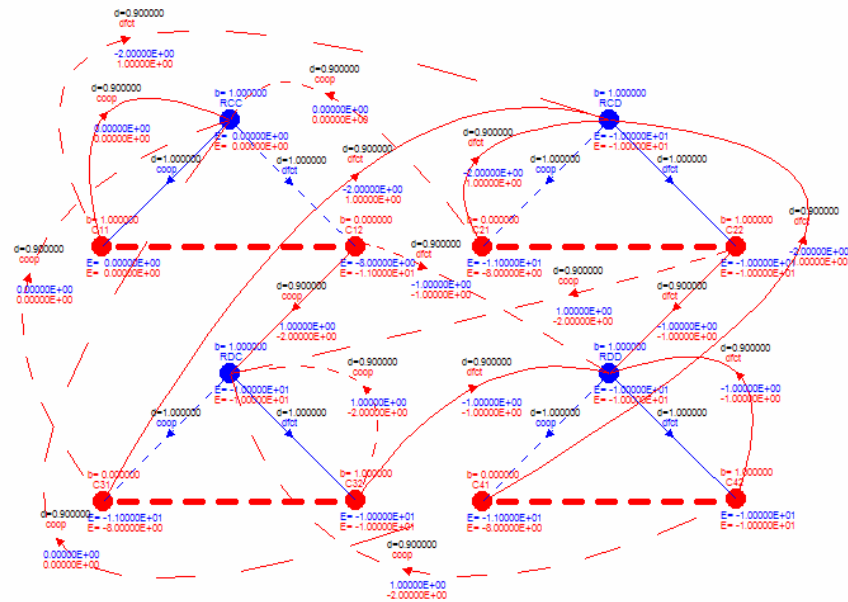


Figure 16: The Grim Trigger in the Repeated Prisoner’s Dilemma

Blue Node RCC is reached anytime both players chose “coop” (for cooperate) on the previous turn. Similarly RCD is reached when blue cooperates and red defects (“dfct”), RDC is reached when blue defects and red cooperates, and RDD results from defection

by both. All red moves have a transient payoff as in Table 7 and bear a $d=0.9$ discount factor. All blue moves have a $d=1$ discount since discounting occurs only when moving to the next turn, and that results from a red move. All dotted moves have probability zero and solid ones therefore have probability one. The SPE of Figure 16 is known as the “Grim Trigger”: any defection by either side leads to perpetual defection. But simultaneous cooperation, if unbroken, will perpetuate itself in the upper left loop.

In general, cooperation can be sustained by trigger schemes provided the discount factor is close enough to 1. The higher the discount factor, the more relevant the future is in rational calculations. This effect is known as the “shadow of the future”.

Homework

1. Explain why the Nash equilibrium of Figure 8 is not a PBE (hint: express the beliefs of the green player and argue that sequential rationality fails).
2. Change the prior beliefs to $P(\text{soft})=0.6$ in Figure 12, solve using *GamePlan* (with the Perfect|Explore solve option), and justify the posterior beliefs in each solution.
3. A game of “perfect information” is one that lacks non-trivial information set. Explain why SPE and PBE must be equivalent in such games.
4. Modify the game of Figure 5 so that the repeated “bid” moves are discounted by factor $d=0.99$ and adjust the repeated bid payoffs as in Figure 15. Solve the resulting game with *GamePlan* and interpret the fact that the probability of bidding increases to $p=0.808081$ (from $p=0.8$) as a result.
5. Solve the game of Figure 16 in Explore|Nash mode and identify “Tit-for-tat” among the solutions. According to the biblical Tit-for-tat, one should retaliate in kind for defection and reciprocate cooperation. Why is that not a SPE?
6. Adjust all discount factors on the red moves to $d=0.49$ in the game of Figure 16. Solve with *GamePlan* in Pure|Explore mode. Do you still find the Grim Trigger?

Brief References

Aumann, Robert, "Survey of Repeated Games," in *Essays in Game Theory and Mathematical Economics in Honor of Oskar Morgenstern*, Vol 4 of Gesellschaft, Recht, Wirtschaft, Wissenschaftsverlag, edited by V. Bohm (1981).

Cournot, Augustin, *Researches into the Mathematical Principles of Wealth* (1838).

Harsanyi, John, “Games with Incomplete Information Played by Bayesian Players, Parts I, II, III”, *Management Science* (1967).

Nash, John, "Equilibrium Points in N-person Games,," *Proceedings of the US Academy of Sciences* (1950).

Selten, Reinhardt, "Re-examination of the Perfectness Concept for Equilibrium Points in Extensive games," *International Journal of Game Theory* (1975).

von Neumann, John, “Zur Theorie der Gesellschaftsspiele (Theory of Parlor Games),” *Mathematische Annalen* (1928).

The GamePlan Software

Please see the *GamePlan* manual. The software can be downloaded free at <http://www.gametheorysoftware.net>

Examples of Game Modeling

Please see Jean-Pierre Langlois' home page on Game Theory.