Strengthening the Weak Perfect Bayesian Solution Concept

Definition 62 *(Kreps and Wilson) A WPBNE (σ,μ) is a* ***sequential equilibrium*** *if there exists a sequence of completely mixed strategies $\left(\sigma^k\right)_{k=0}^{\infty}$ such that*

$$\lim_{k\to\infty}\sigma^k=\sigma$$

and

$$\lim_{k\to\infty}\mu^k=\mu$$

where $\left(\mu^k\right)_{k=0}^{\infty}$ denotes the beliefs derived from $\left(\sigma^k\right)_{k=0}^{\infty}$ using Bayes Rule.

In other words, we introduce completely mixed strategies so that Bayes Rule surely applies.

Example 63 *9.C.4 Off the equilibrium path beliefs – should they be "sensible" in some way? Any sequential equilibrium assigns equal probability to each node in each agent's information set.*

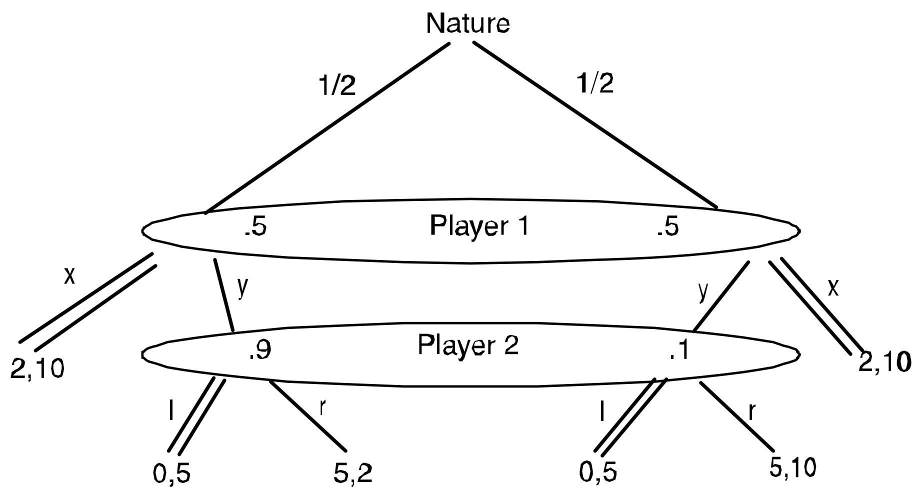

Supppose Player 1 plays y with probability $\varepsilon>0$. Then Player 2 must assign equal probability to each node. Consequently, the above equilibrium can't be sequential.

In a sequential equilibrium, 2 must play r and 1 must play y, with probability .5 assigned to each node in each information set. The above argument determines Player 2's beliefs at his information set. Once these beliefs are determined, it is easy to complete the equilibrium by determining the players' choices.

Example 64 *9.C.5 Returning to the example above in which the WPBNE need not be subgame perfect.*

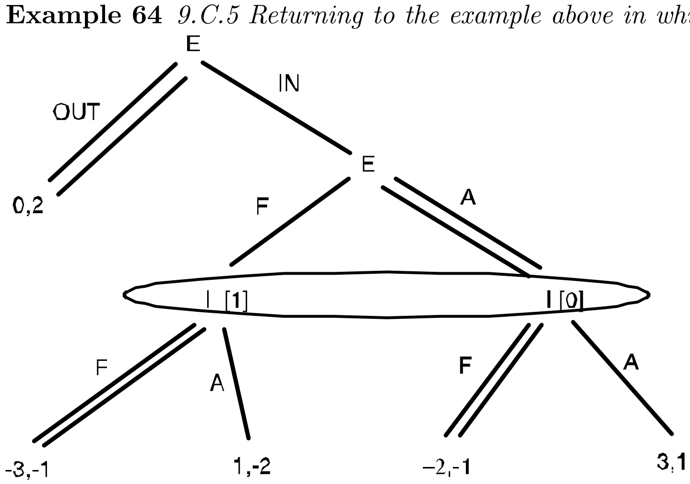

Suppose E chooses In with probability $\delta > 0$ and F with probability $\varepsilon > 0$. Then I must assign probability ε to the left hand node and $1 - \varepsilon$ to the right hand node. In the limit, we have I assigning probability 1 to the right hand node. The above equilibrium is thus not a sequential equilibrium.

Sequential equilibrium is enough to insure that I's beliefs must be consistent with E's strategy after entry. As above,

I chooses A over F if $\mu \leq \frac{4}{5}$.

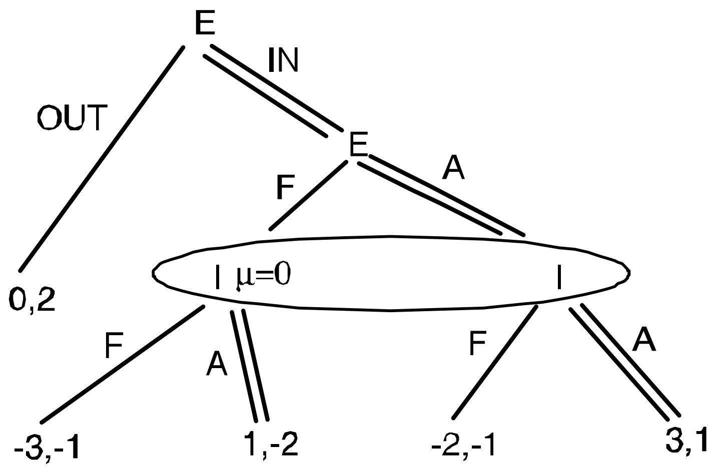

E will therefore choose In and A, and the equilibrium is completed by setting $\mu = 0$.

Unique sequential eq:

E: (In, Accomodate if in)

I: (Accomodate if E plays "in")

Notice that a sequential equilibrium is necessarily a SPNE:

Proposition 65 *A sequential eq. is necessarily a subgame perfect Nash equilibrium.*

Thus, sequential equilibrium strengthens both subgame perfection and weak perfect Bayesian Nash equilibriu

Behavioral motivation for sequential equilibrium? It seems to work, but why is it the right way to refine WPBNE?

9.D. Reasonable Beliefs and Forward Induction

Backward induction (and subgame perfection) models a person who anticipates *future* rational consequences to his actions (i.e., evaluating his choices, he presumes best responses in the future by himself and his opponents, or Nash equilibrium). *Forward induction* concerns the sensibleness of a player's actions and beliefs based upon the *preceding* moves in the game (i.e., a player reasons about what could rationally have happened in the past). This notion is ill-formulated; we aren't going to end up with a definition here that wraps everything up nicely. Instead, we'll identify a set of problems or puzzles that game theorists are still trying to resolve. In this sense, it's a bit like the centipede game: we may have an answer for the game, but it just doesn't seem to be right.

Forward induction addresses the assumption of knowing another player's strategy and knowing that he will stick to it throughout the play of the game. How do players come to know each other's strategy before the game starts? This doesn't model the way play unfolds over time in many situations. This has bothered many of you so far, but it has proven difficult to formalize a theory of games without this assumption. It is the assumption that distinguish3es Nash equilibrium over and above rationalizability, namely, correctness of beliefs about opponents' choices.

The issue in both of the following examples is off the equilibrium path beliefs, namely I assigning positive probability to E playing a strictly dominated strategy off the equilibrium path.

Example 66 *9.D.1 a This is a weak perfect Bayesian equilibrium. In fact, it is a sequential equilibrium. But consider Firm I's beliefs.*

If Firm I finds himself at the information set, he is "certain" that he is at the left hand node. This justifies the choice of F. The strategy In2, however, strictly dominates In1 for E. If Firm I finds himself in the position of having to move, shouldn't he presume that E has put him in that situation by choosing In2? If he does, of course, then he would choose A, in which case E does in fact choose In2.

Notice that there is another sequential equilibrium: E chooses In2, Firm I is certain that he's at the right hand node and chooses A. Forward induction thus serves as a principle that helps us to select one sequential equilibrium as more believable than the other.

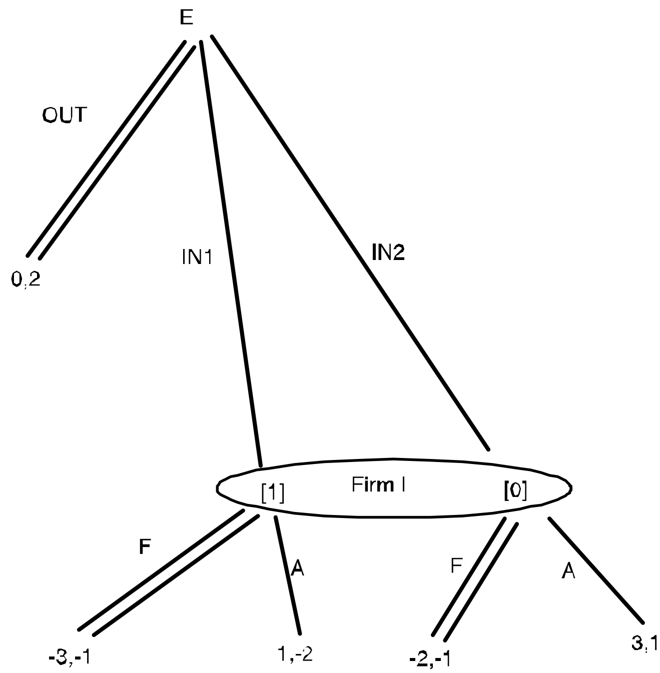

9.D.1b. The Incumbent I here chooses "large niche" because he is certain that he is at the left hand node. This causes the entrant to choose OUT.

Here, "small niche" is strictly dominated for E by OUT. Suppose I finds himself at his information set and has to move. Why would he believe that E chose "small niche"? "Large niche" is not dominated for E by OUT, and so I might conclude that E choose "large niche" in the hope or expectation that I would then choose "small niche".

We have an alternative sequential equilibrium: E chooses "large niche", and I, knowing he's at the right hand node, chooses "small niche". Forward induction perhaps serves as a refinement that helps us to choose this second sequential equilibrium as more reasonable or believable to the outside observer. Firm I acts assuming that E acted rationally to start the game.

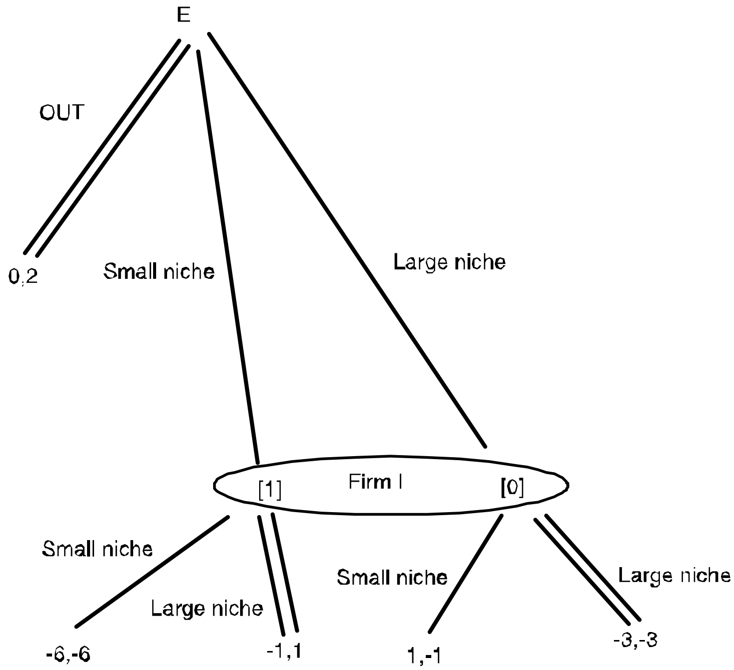

With forward induction, a player chooses a move based upon his analysis of the preceding play of the game, and his assumption that all prior moves have been rational. Forward induction may not be so convincing, however, if there is simply a possibility that players make mistakes.

Text: **"Clearly, the issues here, although interesting and important, are also tricky."**

A Digression on Risk Aversion

Question: We've been dealing with mixed strategies and beliefs in games. How do we incorporate risk aversion?

Let's begin by backtracking a moment. If there is no uncertainty in the game, then the possible outcomes needn't have numerical utility values for the players. We can analyze choices simply by assuming that the players have preferences over the outcomes. This is useful (for instance) in political examples in which the outcomes are selections of candidates.

Example 67 *Two voters (1 and 2) will choose among 3 alternatives (A, B, and C) by successively eliminating choices: 1 goes first and eliminates A, B, or C, and 2 then eliminates one of the two remaining choices, which determines the winning choice. We can draw this as an extensive form game, and we can analyze the game by backwards induction if each player has complete and transitive preferences over the choices (e.g., A>B>C). Backwards induction can be complicated by multiple equilibria if either player is indifferent among some outcomes, but strict preferences are not required to analyze the game by backwards induction.*

We need numerical values of utility, however, if we are to calculate expected return. Mixed strategies and nontrivial beliefs at information sets require a utility representation of preferences. Risk aversion is incorporated by the utility assignments to outcomes (specifically, the utility received from money).

When we consider a particular game with numerical payoffs given for the various outcomes, risk aversion has already been incorporated in the numbers assigned to the outcomes. Changing risk preferences, or examining the effect of risk aversion, requires changing the payoffs associated with the outcomes of the game.

Throughout our discussion of game theory, we have assumed that players know the structure of the game that they are playing. If there is complete information, then this effectively means that each player knows the "risk preferences" of his opponents. That is, if there are monetary payoffs in the game and complete information, then each player knows the utility that any other player assigns to the different monetary outcomes of the game. This is a big assumption!

Appendix A: Finite and Infinite Horizon Bilateral Bargaining

Alternating offers bargaining, with acceptance ending the game. $v = 1$ dollars to be split (MWG allows v arbitrary, but little is gained from this generality).

Discount factor $\delta_i < 1$ of player i; starting in period 1, a dollar received in period t is worth δ_i^{t-1}. This puts a "time pressure" on the players to reach agreement.

- Why are economists interested in bilateral bargaining?

  It is a foundational or fundamental economic relationship on which the theory of markets should be built.

- Why are we interested in infinite horizon bargaining?

  It models bargaining without a predetermined deadline or ending point known to the bargainers. It is also the case that we tend to get trivial or less meaningful results in finite models.

- How do we interpret a discount factor δ_i?

  Either as the time cost of money or the probability of continuing to another stage (if all players share the same discount factor).

- A model of alternating offers bargaining that was formulated separately by Dale Stahl and Ariel Rubenstein.

Example 68 *Finite horizon: bargaining for 2 stages. First, recall the 1-stage analysis. Notice that this is not a repeated game, as there is only one dollar to be divided over time (not a new one for each stage). For convenience, we'll always have the offer made in terms of 1's share (even when 2 makes the offer).*

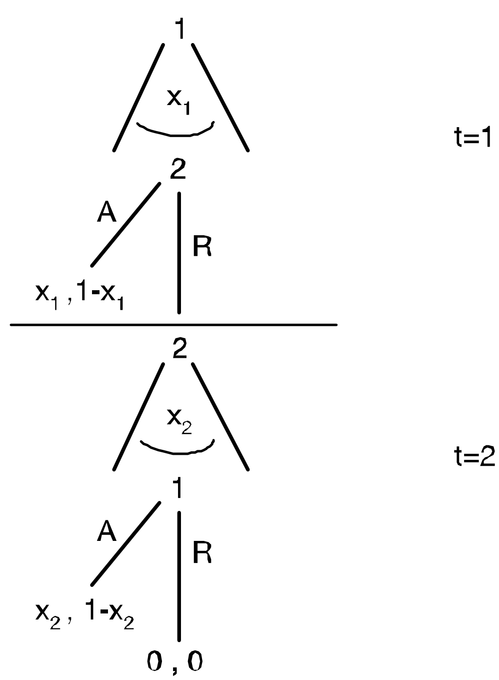

- *Any division of the dollar in either stage 1 or stage 2 can be sustained by a Nash equilibrium. Explain why these equilibria are not subgame perfect.*
- *Subgame perfection determines a unique division of the dollar. In fact, there is a unique SPNE.*
- *Nash equilibrium has two useful implications that are worth stating:*

  1. *If i makes j an offer that j will accept, then the best response property of a Nash equilibrium requires that* ***i offer j the minimal amount that j will accept****, according to j's strategy.*

2. *Player j's rule for which offers he accepts at a given stage is based upon the discounted value in that stage of rejecting the offer: j should reject an offer that gives him y in stage t if and only if the present value in stage t of what he'll get by rejecting exceeds y. The present value of continuing to the next stage thus determines a player's cutoff for accepting/rejecting offers.*

Example 69 *Solve in the case of the $T = 3$ stage game. Do it once using the solution for the $T = 2$ game.*

Proposition 70 *There is a unique SPNE in the T-stage game, where T is finite.*
We've shown how to solve for it, by backwards induction.

Example 71 *Infinite Horizon. Backwards induction isn't helpful here because there isn't a "bottom" of the game to start at. We will use the symmetry and self-similarity of subgames of the games to deduce the answer.*

Let M denote the maximum payoff that player 1 receives in any SPNE and let m denote the lowest amount player 1 receives in any SPNE. (More formally, we should use "sup" (or "least upper bound") and "inf" (or "greatest lower bound").) We're assuming for now that SPNE exist, which we'll show by construction in the analysis.

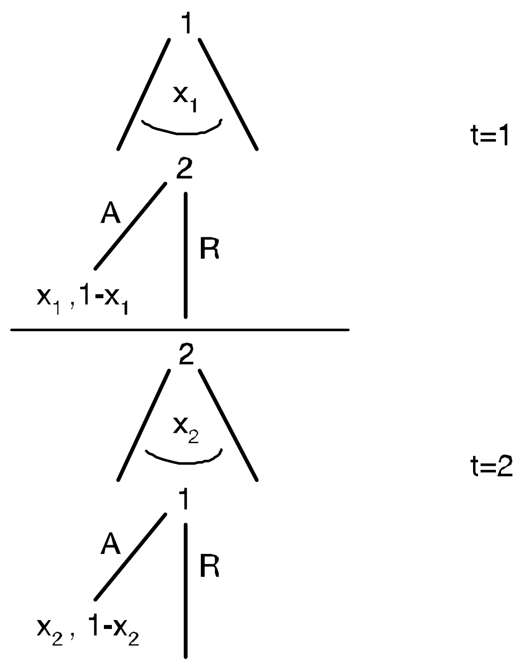

Using m and M, we'll compute other lower and upper bounds on SPNE payoffs from subgame perfection and the structure of the game.

Step 1. ***Determination of SPNE payoffs.*** *The table below lists payoffs in the dollar value of the stage in which they are listed.*

subgame starting at:	*proposal made by*	*1 gets at least*	*2 gets at most*
stage 1	*1*	$(++) 1 - \delta_2 (1 - \delta_1 m)$	
stage 2	*2*	$(**)\, \delta_1 m$	$(+)\, 1 - \delta_1 m$
stage 3	*1*	$(*)\, m$	

- *The entry $(*)$ represent the fact that the subgame starting in stage 3 is identical to the entire game.*
- *Since 1 can get at least m in stage 3, he should reject any offer in stage 2 that doesn't give him at least $\delta_1 m$. This explains the entry $(**)$.*
- *If 1 gets at least $\delta_1 m$, then 2 gets at most $1 - \delta_1 m$, since they only have one dollar to divide. This explains $(+)$.*

- *Finally, 2 can get at most $\delta_2(1-\delta_1 m)$ in stage 1 dollars by declining 1's offer in stage 1 and going on to stage 2. 2 would therefore accept an offer from 1 that gave him $\delta_2(1-\delta_1 m)$, and so 1 can get at least $1-\delta_2(1-\delta_1 m)$ by setting x_1 equal to this value (which 2 then accepts). It therefore must be the case that 1 gets at least $1-\delta_2(1-\delta_1 m)$ in a SPNE, which implies the value (++): 1 gets at least m in a SPNE, which is therefore at least $1-\delta_2(1-\delta_1 m)$ (a lower bound we've derived on what he gets in a SPNE).*

We therefore have

$$1-\delta_2(1-\delta_1 m) \leq m$$

or

$$1-\delta_2 \leq (1-\delta_2\delta_1)\, m \Leftrightarrow$$
$$\frac{1-\delta_2}{1-\delta_2\delta_1} \leq m.$$

Step 2. *Determination of SPNE payoffs (cont.).* *Similarly, consider the following table:*

subgame starting at:	*proposal made by*	*1 gets at most*	*2 gets at least*
stage 1	*1*	$1-\delta_2(1-\delta_1 M)$	$\delta_2(1-\delta_1 M)$
stage 2	*2*	$(+)\ \delta_1 M$	$(**)\, 1-\delta_1 M$
stage 3	*1*	$(*)\, M$	

- *The entry $(*)$ represents the fact that the subgame starting in stage 3 is identical to the entire game.*
- *Since 1 can get at most M in stage 3, he should accept any offer in stage 2 that gives him at least $\delta_1 M$. This explains the entry $(**)$ (if 2 made this offer in stage 2, then 1 would accept it).*
- *If 2 gets at least $1-\delta_1 M$, then 1 gets at most $\delta_1 M$, since they only have one dollar to divide. This explains $(+)$.*
- *Finally, 2 gets at least $\delta_2(1-\delta_1 M)$ in stage 1 dollars by declining 1's offer in stage 1 and going on to stage 2. 2 would therefore reject any offer from 1 that gave him less than $\delta_2(1-\delta_1 M)$, and so 1 can get at most the maximum of*
$$\delta_2\delta_1 M \text{ and } 1-\delta_2(1-\delta_1 M)$$
in stage 1 dollars. The second term is what he gets if the offer of $x_1 = \delta_2(1-\delta_1 M)$ is accepted, while the first is the value in stage 1 dollars if this offer is rejected. It is clear that
$$1-\delta_2(1-\delta_1 M) = 1-\delta_2+\delta_2\delta_1 M > \delta_2\delta_1 M,$$
and so $1-\delta_2(1-\delta_1 M)$ is the bigger of these two numbers. Consequently, we have found that the maximum M of what 1 gets in a SPNE is bounded above by $1-\delta_2(1-\delta_1 M)$.

We therefore have

$$M \leq 1-\delta_2(1-\delta_1 M)$$

or

$$(1-\delta_2\delta_1)\, M \leq 1-\delta_2 \Leftrightarrow$$
$$M \leq \frac{1-\delta_2}{1-\delta_2\delta_1}.$$

Putting our two inequalities together, we have

$$M \leq \frac{1-\delta_2}{1-\delta_2\delta_1} \leq m,$$

but $m \leq M$ by definition. We've proven that there is only one possible payoff for 1 in a SPNE:

$$\frac{1-\delta_2}{1-\delta_2\delta_1}.$$

Step 3. *Determination of SPNE strategies:* *The subgame starting in stage 2 is just like the entire game, only with the roles of the players reversed. We thus can fill in the entries (*) below:*

subgame starting at:	*proposal made by*	*1 gets*	*2 gets*
stage 1	*1*	(*) $\frac{1-\delta_2}{1-\delta_2\delta_1}$	
stage 2	*2*		$(*)\ \frac{1-\delta_1}{1-\delta_2\delta_1}$

- *2 will accept x_1 in stage 1 if and only if*

$$1 - x_1 \geq \delta_2 \left(\frac{1-\delta_1}{1-\delta_2\delta_1} \right)$$

or

$$\begin{aligned} 1 - \delta_2 \left(\frac{1-\delta_1}{1-\delta_2\delta_1} \right) &\geq x_1 \\ \frac{1-\delta_2\delta_1 - (\delta_2 - \delta_2\delta_1)}{1-\delta_2\delta_1} &\geq x_1 \\ \frac{1-\delta_2}{1-\delta_2\delta_1} &\geq x_1 \end{aligned}$$

By making this offer, 1 and 2 both receive what we have determined as their SPNE payoffs in the game. If 1 instead chose a lower x_1, the offer would be accepted but 1 wouldn't receive his SPNE payoff. If 1 chose a larger x_1, then the game would go on to the next stage. 1's payoff in stage 1 dollars would therefore be at most

$$\begin{aligned} \delta_1 \left(1 - \frac{1-\delta_1}{1-\delta_2\delta_1} \right) &= \delta_1 \left(1 - \frac{1-\delta_1}{1-\delta_2\delta_1} \right) \\ &= \delta_1 \left(\frac{1-\delta_2\delta_1 - (1-\delta_1)}{1-\delta_2\delta_1} \right) \\ &= \delta_1 \left(\frac{-\delta_2\delta_1 - (-\delta_1)}{1-\delta_2\delta_1} \right) \\ &= \delta_1^2 \left(\frac{1-\delta_2}{1-\delta_2\delta_1} \right) \end{aligned}$$

which is less than his SPNE payoff. The offer must therefore be accepted. We have thus determined that 1 offers

$$x_1 = \frac{1-\delta_2}{1-\delta_2\delta_1}$$

while 2 accepts an offer of x_1 if and only if

$$1 - x_1 \geq \delta_2 \left(\frac{1-\delta_1}{1-\delta_2\delta_1} \right)$$

Using the symmetry of the game, and the fact that subgames have the same structure as the entire game, it is easy to deduce the offers that would be made in each period along with the responder's rule for accepting/rejecting.

- *The SPNE division of the dollar is therefore*

$$x_1 = \frac{1-\delta_2}{1-\delta_2\delta_1}$$

Notice that x_1 is increasing in δ_1 and decreasing in δ_2. This is reasonable if we interpret δ_i as a measure of player i's "patience", or ability to endure delay before reaching agreement.

- *Equilibrium strategies: Player 1 always offers*

$$x_i = \frac{1-\delta_2}{1-\delta_2\delta_1}$$

whenever he makes the proposal. He accepts the proposal of x_i by 2 only if

$$x_i \geq \delta_1 \left(\frac{1-\delta_2}{1-\delta_2\delta_1}\right),$$

i.e., the present value of what he will get tomorrow when he again makes the offer. By symmetry, Player 2 always proposes

$$1 - x_i = \frac{1-\delta_1}{1-\delta_2\delta_1} \Leftrightarrow$$
$$x_i = 1 - \frac{1-\delta_1}{1-\delta_2\delta_1} = \delta_1 \left(\frac{1-\delta_2}{1-\delta_2\delta_1}\right)$$

whenever he makes the proposal. He accepts an offer of x_i from Player 1 only if

$$1 - x_i \geq \delta_2 \left(\frac{1-\delta_1}{1-\delta_2\delta_1}\right) \Leftrightarrow$$
$$x_i \leq 1 - \delta_2 \left(\frac{1-\delta_1}{1-\delta_2\delta_1}\right) = \frac{1-\delta_2}{1-\delta_2\delta_1}.$$

This defines the strategies completely, i.e., in every stage and both on and off the equilibrium path.

Summary of Alternating Offers Bargaining

1. Any finite version of the game has a unique SPNE that is obtainable by backwards induction. Any division of the dollar in any stage, however, is sustainable in a Nash equilibrium.

2. Any division of the dollar in any stage is sustainable in a Nash equilibrium of the infinite game. The infinite game has a unique SPNE that of course is not derived by backwards induction. It is instead derived using the fact that any subgame of the infinite game is identical in structure to the entire game (perhaps with the roles of the two players interchanged), and so any statement derived for a SPNE in the entire game holds in every subgame.

3. Derivation of the SPNE in the infinite game:

   (a) Let m denote the minimum payoff to player 1 in any SPNE and M the maximum payoff in any SPNE. We derive an lower bound on m and an upper bound on M. These bounds are the same ($=\frac{1-\delta_2}{1-\delta_2\delta_1}$), however, and $M \geq m$, which allows us to conclude that this bound is the amount that player 1 obtains in any SPNE.

   (b) The game starting in stage 2 looks just like the entire game with the roles of the players reversed. This allows us to conclude that the value of being in the game in stage 2 to player 2 is $\frac{1-\delta_1}{1-\delta_2\delta_1}$.

   (c) Using 3.b., we then deduce player 2's cutoff for accepting/rejecting offers in stage 1 (A iff he gets at least $\delta_2 \cdot \frac{1-\delta_1}{1-\delta_2\delta_1}$). This in fact defines his rule for accepting/rejecting in any stage.

   (d) Using 3.c, we can deduce player 1's cutoff in any stage in which he responds to an offer by player 2 (A iff he gets at least $\delta_1 \cdot \frac{1-\delta_2}{1-\delta_2\delta_1}$). We could also conclude this from 3.a.

   (e) If player 1 offers player 2 $\delta_2 \cdot \frac{1-\delta_1}{1-\delta_2\delta_1}$ to start the game, then player 2 accepts. Player 1 then gets

   $$1 - \delta_2 \cdot \frac{1-\delta_1}{1-\delta_2\delta_1} = \frac{1-\delta_2\delta_1-\delta_2+\delta_2\delta_1}{1-\delta_2\delta_1} = \frac{1-\delta_2}{1-\delta_2\delta_1},$$

   which is his payoff in any SPNE of the game. If he offers player 2 less than this amount, then player 1 gets at most $\delta_1 \cdot \delta_1 \cdot \frac{1-\delta_2}{1-\delta_2\delta_1}$, which is the present value of the most that he can get by going to the next stage. We have

   $$\frac{1-\delta_2}{1-\delta_2\delta_1} > \delta_1 \cdot \delta_1 \cdot \frac{1-\delta_2}{1-\delta_2\delta_1},$$

and so player 1 prefers to make an acceptable offer in stage 1. We thus deduce that whenever 1 makes an offer, it is

$$x_t = \delta_2 \cdot \frac{1-\delta_1}{1-\delta_2\delta_1}.$$

(f) Using 3.e, we deduce that whenever player 2 makes an offer, he proposes giving player 1

$$\delta_1 \cdot \frac{1-\delta_2}{1-\delta_2\delta_1}.$$

In the notation of the last class, this means

$$x_t = 1 - \delta_1 \cdot \frac{1-\delta_2}{1-\delta_2\delta_1}.$$