Nash Equilibrium:

Self-Enforcing Agreement, Self-Fulfilling Beliefs, Randomization, Evolution

Adam Brandenburger*

Version 01/04/07

Theorem 1 in the note "Nash Equilibrium: Definition" gives conditions on the players' rationality and beliefs that yield Nash equilibrium. This leaves open the question of how such a configuration of beliefs might come about. This note, which is very informal, looks at this and some other questions about Nash equilibrium.

1 Self-Enforcing Agreement

A common–perhaps the most common–way of thinking about Nash equilibrium is as a self-enforcing agreement. More precisely, the argument goes that if the players reach an agreement about how to play the game, then a necessary condition for the agreement to hold up is that the strategies constitute a Nash equilibrium. Otherwise, it is argued, at least one player will have an incentive to choose a strategy different from the agreed-upon one.

One (not necessarily compelling) objection that can be made to this argument is that it envisages a more or less explicit process of preplay communication among the players. If this takes place, then perhaps it ought to be incorporated into the description of the game. That is, we should be dealing with a larger game, in which these communication moves are explicitly modelled. But then we would be back to square one! (But, one might also decide that communication should *not* be explicitly modelled in the same way as the subsequent moves in the 'real' game.) Bob Aumann raises another difficulty.[1]

Example 1 *Consider the game in Figure 1 below.*

*With the assistance of Amanda Friedenberg. nashmisc-01-04-07

[1] "Nash Equilibria Are Not Self-Enforcing," in *Economic Decision Making: Games, Econometrics, and Optimisation: Essays in Honor of Jacques Dreze*, edited by J. J. Gabszewicz, J.-F. Richard, and L. Wolsey, pp.201-206, Elsevier Science Publishers, Amsterdam, 1990.

		Player II	
		Current	*New*
Player I	*Current*	9, 9	0, 8
	New	8, 0	7, 7

Figure 1

This game is called the Security Dilemma in the international relations literature.[2] *Two countries can each choose to stay with the current defense system, or invest in a new and expensive system. Each country would prefer to stay with the current system if the other country makes the same decision, but would become very vulnerable if the other country invested in the new system and it didn't.*

There are two pure-strategy Nash equilibria (Current, Current) and (New, New). (There is also one mixed-strategy Nash equilibrium, as you should check.) Aumann argues that (Current, Current) cannot be considered a self-enforcing agreement. Suppose player I tries to reach an agreement with player II that both will choose Current. The difficulty is that I is happy to have II choose Current even if he (I) intends to play New, since this raises his payoff from New to 8 *from* 7. *Player I is motivated to make the agreement, regardless of what he intends to play. Therefore, Aumann argues, any agreement between the players to play (Current, Current) is useless.*

Note carefully that Aumann isn't saying that (Current, Current) can't happen, just that an agreement to play this has no effect.

2 Self-Fulfilling Beliefs

Another view of Nash equilibrium–which we already mentioned in "Nash Equilibrium: Definition"–is that it is the outcome of a learning process. Suppose players play a game repeatedly. The idea is that a reasonable learning process, if it converges, should converge to a Nash equilibrium. The argument goes that if the strategies chosen do indeed settle down (i.e. converge), then each player will come to assign probability 1 to the actual strategies chosen by the other players. In this case, the conditions of Theorem 1 in "Nash Equilibrium: Definition" apply, and the strategies must form a Nash equilibrium.

[2] "Cooperation under the Security Dilemma," by R. Jervis, *World Politics*, 30, 1978, 167-214.

Two immediate comments: First, this approach appears to presuppose some kind of bounded rationality on the part of players. Otherwise, they will view the repeated game as a single game, and then we're back to square one. (Of course, the bounded rationality assumption may be quite reasonable.) Second, the argument is a conditional one–conditional on there being convergence. Non-convergence is, presumably, a real possibility.

There is also another, more subtle issue. The argument that convergence, if takes place, must be to a Nash equilibrium turns out to be quite incomplete. The issue is best seen via an example.[3]

Example 2 *Consider the game in Figure 2.*

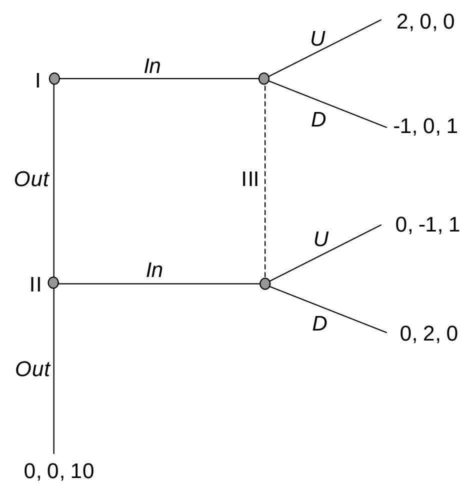

Figure 2

Suppose that player I believes that player III's strategy is D, while player II believes that player III's strategy is U. Then it seems reasonable to conclude that player I will choose the strategy Out and player II the strategy Out. But there is no (pure-strategy) Nash equilibrium in which Out and Out are played. To see this, suppose that player III chooses U in the equilibrium. Then player I will choose In. Or, if player III chooses D in the equilibrium, player II will choose In. (There is a little more to check here. Can you see what it is?)

The point is that while the situation in which Out and Out are played is not a Nash equilibrium, it is 'self-fulfilling' in the following way. The game is played over and over. Player I chooses Out and player II chooses Out. Since

[3] Similar to ones in the book *The Theory of Learning in Games,* by Drew Fudenberg and David Levine (MIT Press 1998). Note that the example involves a game tree, an object which we haven't formally introduced yet. In fact, the whole discussion of the example is a heuristic one, and should be treated as such.

player III never gets to move (neither I nor II chooses In), neither player I nor player II will ever find his hypothesis about player III refuted. Players I and II can reasonably maintain their respective beliefs about player III's strategy–and hence go on playing Out and Out. True, the two players disagree about player III, but neither will have reason to change his mind.

(A counter to this is that either player I or player II might decide to test his hypothesis by giving player III the move–i.e. by choosing In. And then there will be an observation on player III's strategy.)

3 Randomization

Apart from the question of how Nash equilibrium might come about, the presence of mixed strategies in the definition of equilibrium also raises a question of interpretation.

The use of mixed strategies has always been conceptually troubling. For one thing, there is never a strict incentive for a player to randomize. To see this, go back to the game in Figure 2 in "Nash Equilibrium: Definition." Given player 2's Nash-equilibrium strategy (probabilities $\frac{1}{3}$ on *Left* and $\frac{2}{3}$ on *Right*), player 1 is indifferent between carrying out the $\frac{1}{2}:\frac{1}{2}$ randomization given by the equilibrium and any other randomization between *Up* and *Down* (including choosing either *Up* or *Down* for sure). This would not matter if no such 'deviation' by player 1 upset the equilibrium, but this is not the case. It is only when player 1 performs the $\frac{1}{2}:\frac{1}{2}$ randomization that player 2 is prepared to play $\frac{1}{3}:\frac{2}{3}$. Apart from this difficulty, there is the feeling that perhaps people do not really randomize when making decisions.

One response is to think in terms of an evolutionary model, as discussed in the next section. Another response is simply to do away with the assumption that players randomize, and suppose instead that each player chooses some definite pure strategy–exactly where we began. The key idea is that the other players need not know which one. Each player's mixed strategy can then be thought of as representing not a conscious randomization by that player, but rather the common probabilistic assessment made by the other players about that player's choice. It is *randomization as ignorance*, if you like.

The general approach in this vein is to imagine each player to have beliefs about the game (including other players' strategy choices), about the other players' beliefs about the game, etc. This is what in the literature is called the *epistemic* approach to game theory. For details on how mixed-strategy Nash equilibrium can be understood in these terms, refer to "Epistemic Conditions for Nash Equilibrium," by Robert Aumann and Adam Brandenburger, *Econometrica*, 1995, 63, 1161-1180.

4 Evolution

Example 3 *Consider the game in Figure 3. Here, two animals are interested in a certain prize (such as a piece of territory). Think of the prize as worth* 2 *units, and suppose that fighting costs* 3 *units. If both animals fight (choose the Hawk strategy), then each receives the prize with probability* $\frac{1}{2}$*. Thus, each has an expected payoff of* -2*. If one animal does not fight (chooses the Dove strategy), while the other chooses the Hawk strategy, then the latter animal obtains the prize of 2. Finally, if both animals choose the Dove strategy, they share the prize.*

1 \ 2	*Hawk*	*Dove*
Hawk	-2, -2	2, 0
Dove	0, 2	1, 1

Figure 3

There are two pure-strategy Nash equilibria of the game: One animal chooses *Hawk* and the other chooses *Dove*. There is also a mixed-strategy equilibrium in which each animal chooses *Hawk* with probability $\frac{1}{3}$ and *Dove* with probability $\frac{2}{3}$. (To verify that this is an equilibrium, note that the strategy *Hawk* then yields an expected payoff of $\frac{2}{3}$, which is the same as that yielded by the strategy *Dove*.)

So far, there is nothing new here. But there is another way to think about this game, in which we imagine that there is a (large) population of animals and that members of this population encounter one another at random. Mixed strategies can then be thought of as distributions of strategies in the population. Thus, the mixed strategy above corresponds to the case in which $\frac{1}{3}$ of the population employs the *Hawk* strategy while $\frac{2}{3}$ of the population employs the *Dove* strategy. Moreover, this proportion appears to have a kind of stability property. If the proportion of Hawks were $\frac{1}{2}$, say, then Doves would have higher expected payoffs than Hawks, and, for this reason, might be expected to increase in the population. Likewise, a population consisting exclusively of Hawks, or exclusively of Doves, would not be stable.

This type of analysis belongs to a very interesting sub-field of game theory called Evolutionary Game Theory.