Nash Equilibrium

Larry Blume

Table of Contents

I Introduction

Suppose that a model of a game includes all relevant information needed for forecasting rational choices in a game. Suppose too that there is a *knowable*, unique way to play each game. Then that theory of playing a game must forecast a Nash equilibrium. Although Nash equilibrium is not without its problems, it is **the** basic solution concept, and the subject of movies.

II Definitions and Theorems

Nash equilibrium is a *normal form* solution concept. Let $G = \big(I, (S_i, u_i)_{i \in I}\big)$ denote a game in normal form, not necessarily finite.

Definition 1. *A strategy profile $s^* \in S$ is a Nash Equilibrium of the game G iff for all $i \in I$, $u_i(s_i^*, s_{-i}^*) \geq u_i(t_i, s_{-i}^*)$ for all $t_i \in S_i$.*

The existence of Nash equilibrium is not guaranteed without some assumptions on preferences, strategy sets, and so forth. First we take up Nash equilibrium in finite games. Not all finite games have a Nash equilibrium. Matching pennies, for instance, does not. But Nash equilibria exist for the mixed extensions of all finite games.

Definition 2. *A mixed strategy profile σ^* of a finite game G is a mixed strategy Nash equilibrium of G iff it is a Nash equilibrium of the mixed extenstion $\hat{G}$.*

Recall that for a mixed strategy profile σ,

$$U_i(\sigma) = U_i(\sigma_i, \sigma_{-i}) = \sum_s u_i(s_i, s_{-i})\sigma_i(s_i)\sigma_{-i}(s_{-i}),$$

where by $\sigma_{-i}(s_{-i})$ I mean the probability $\prod_{j \neq i} \sigma_j(s_j)$.

Suppose now that $\hat{G}$ is the mixed extension of some finite game $G = (I, (S_i, u_i)_{i \in I}$. Here is an equivalent definition of Nash equilibrium.

Definition 3. *A strategy profile* σ^* *is a mixed strategy Nash Equilibrium of the game* $\hat{G}$ *iff for all* $i \in I$ *and* $s_i \in S_i$*, if* $\sigma_i^*(s_i) > 0$ *then*

$$U_i(s_i, \sigma_{-i}^*) \geq U_i(t_i, \sigma_{-i}^*)$$

for all $t_i \in S_i$*.*

You should prove the equivalence. Also the following:

Definition 4. *A strategy profile* σ^* *is a mixed strategy Nash Equilibrium of the game* $\hat{G}$ *iff for all* $i \in I$ *and* $s_i, t_i \in S_i$*, if* $U_i(s_i, \sigma_{-i}^*) < U_i(t_i, \sigma_{-i}^*)$ *then* $\sigma_i^*(s_i) = 0$*.*

John Nash's fundamental contribution to game theory is his statement and proof of the following theorem.

Theorem 1. *Every finite game has a Nash equilibrium in mixed strategies.*

Here is a more general existence theorem that covers infinite games such as the linear Cournot oligopoly model.

Theorem 2. *Let* $G = (I, (S_i, u_i)_{i \in I}$ *denote a game such that:*

1. I *is finite;*

2. *for all* $i \in I$, S_i *is a compact and convex subset of a finite dimensional vector space;*

3. *for all* $i \in I$, u_i *is continuous, and quasi-concave in* s_i *for each* s_{-i}*.*

Then a Nash equilibrium exists.

One can do even better, and this theorem excludes important classes of games, such as Hotelling-style location games and even Cournot oligopoly with a general class of demand functions. But it is all we are going to cover.

III Interpretations

In games such as the Cournot and Bertrand duopoly models it is easy to imagine many kinds of dynamic adjustment by which firms would reach equilibrium quantities and prices. In games like matching pennies, it is hard to imagine playing any other way than by flipping the coins. But in other games Nash equilibrium seems less useful, both for the nature of its predictions and for its use of seemingly exotic mixed strategies. I will take up both of these issues here, albeit incompletely.

III.1 Equilibrium Predictions

Consider the *battle of the sexes*, so-named because it comes with a story from the 50's too non-pc to relate here. This game is emblematic of much social

	B	S
B	$2,1$	$0,0$
S	$0,0$	$1,2$

Figure 1: Battle of the Sexes

conflict because there are gains from cooperation but also conflict over the distribution of the gains. This is an example of a coordination game. This game has three Nash equilibria, two pure and one mixed. The pure equilibria are, of course, B,B and S,S, and the mixed equilibrium has the row player choosing B with probability $2/3$ while the column player chooses B with probability $1/3$. Since the whole point of Nash equilibrium is the conflict between B,B and S,S, saying that both are equilibria is not much more than saying there is a conflict. The mixed equilibrium certainly does not appeal as a prescription of true randomization. In this game it would seem that Nash equilibrium is not of much help, but it is hard to imagine any theory that could make a sharper prediction based on the data available in the normal form.

However one can imagine better theories in the coordination game of Figure 2: Again there are three Nash equilibria: the pure strategy profiles A,A and B,B, and a mixed equilibrium wherein each player choose A with probability $1/3$. The A,A equilibrium is said to be payoff dominant, which is a game theorist's

	A	B
A	$2, 2$	$0, 0$
B	$0, 0$	$1, 1$

Figure 2: Payoff Dominance

way of saying that this equilibrium Pareto-dominates all the other equilibria. Most of us would think that any reasonable theory should predict A, A. It is hard to imagine rational players ending up anywhere else. A "better theory" would predict A, A and none of the other equilibrium. Notice, however, that the Nash property is still salient. Consider the game in Figure 3 Here C, C Pareto dominates A, A,

	A	BC	
A	$2, 2$	$0, 0$	$4, 0$
B	$0, 0$	$1, 1$	$0, 0$
C	$0, 4$	$0, 0$	$3, 3$

Figure 3: Payoff Dominance

and yet we still find A, A to be the unique reasonable prediction. (You should find all the equilibria of this game.) These games suggest that Nash equilibria need be further refined to select among equilibria. Thomas Schelling (1963) introduced the concept of a focal point, a particular equilibrium distinguished in some way that might lead players to play according to its prediction. In these two games the A, A equilibrium is focal because of some feature of the game that is reflected in its normal form. The entire business of equilibrium refinements represents attempts to distinguish equilibria which may be focal due to the payoff structure of the game. Schelling points out that equilibria may be focal for completely inessential reasons, as in the game of Figure 4. Payoff dominance seems

	A	B
A	$6, 6$	$5, 7$
B	$7, 5$	$4, 4$

Figure 4: A Focal Point

to be an attractive principle for a focality requirement internal to the game. The game in Figure 5 may make you rethink this position. If row thinks that there is even a 2% chance that col will choose B, then B is better than A. If this is not

	A	B
A	$2, 2$	$-100, 0$
B	$0, -100$	$1, 1$

Figure 5: Risk Dominance

convincing, replace -100 with -10^6. Naturally a principle has been invented to justify B. Strategy B is said to be risk dominant.

Definition 5. *In a 2×2 coordination game G with pure strategy equilibrium profiles s and t, the profile s is risk dominant iff for all i, $u_i(s) - u_i(t_i, s_j) > u_i(t) - u_i(s_i, t_j)$.*

This is equivalent to saying that for both players, s_i is the unique best reply to any beliefs which assign probability at least $1/2$ to s_j. This can be generalized to larger games and more general classes of games. See Harsanyi and Selten (1988).

III.2 Mixed-Strategy Equilibrium

Mixed strategies pose an interpretive problem, because although we find randomization credible in matching pennies and poker, we don't find it credible in deciding to enter a market, or in setting a tax policy (maybe). Here are four different interpretations of a mixed-strategy equilibrium:

Pure randomization: Players really choose to randomize.

Belief equilibrium: Suppose that σ^* is a mixed-strategy Nash equilibrium. Let $T_i = \{s_i : \sigma_i^*(s_i) > 0\}$. The set $T = \prod_{i \in I} T_i$ is rationalizable. The beliefs that rationalize each $t_i \in T_i$ are those given by σ_{-i}^*.

Population equilibrium: Imagine that each 'player' is in fact a large population of individuals, and players are matched up randomly, one of each type. Then $\sigma_i^*(s_i)$ can be interpreted as the fraction of type i players choosing s_i. In his thesis, Nash refers to this as the mass action interpretation of Nash equilibrium.

Incomplete information: The game theorist John Harsanyi has suggested interpreting mixed-strategy Nash equilibria as the outcome of a Bayesian game, wherein players payoffs are independently perturbed around the original payoff matrix, and each players' payoffs are observable only to him- or herself. We will work this out later.

Mixed strategy Nash equilibria are counter-intuitive. The game in Figure 6 has a unique Nash equilibrium in which the column player plays A with probability

	A	B
A	$3,0$	$0,1$
B	$1,3$	$1,1$

Figure 6: A unique equilibrium which is mixed.

$1/3$ and the row player plays A with probability $2/3$. Suppose now that the 0 payoff for column in the event of A, A is changed to -1. How does the equilibrium change? Although it is column's utility that changes, it is *row's equilibrium strategy* which changes. The game has a unique Nash equilibrium in which row plays A with probability $1/2$, while column remains as before. The equilibrium condition requires that column be indifferent between A and B in order to randomize. If column's payoffs change, it is row's strategy that must change to keep column indifferent. What happens if column's A, A payoff is increased from 0 to $1/2$?

IV Equilibrium and Dominance

Nash equilibrium makes sharper predictions than does iterated elimination of never-best-replies because it imposes further restrictions on what beliefs can be used to justify a given pure strategy:

1. All strategies of player i must be justified by the *same* belief.
2. In player i's beliefs, the choices of players j and k must be independent.
3. Players j and k must share common beliefs about player i.

Notice that the last two points only matter for games with three or more players. Nonetheless it is easy to prove the following theorem (and you should do it).

Theorem 3. *If strategy s_i belonging to player i is played with positive probability in a Nash equilibrium, then it is iteratively undominated.*

The conclusion is not true for weak domination. The game of Figure 7 has two Nash equilibria, A, A and B, B, and the latter uses a weakly dominated strategy. But the first equilibrium does not, and the following theorem is easy to

	A	B
A	$1,1$	$0,0$
B	$0,0$	$0,0$

Figure 7: Equilibrium in weakly dominated strategies.

show (hint — try it):

Theorem 4. *In any finite game G and any path of iterated elimination of weakly domination, there is a mixed-strategy Nash equilibrium σ^* in which for each player i, $\sigma_i^*(s_i) > 0$ implies s_i survives iterated elimination.*

The only issue is to show that strategies which are discarded along the way are not better replies to the equilibrium.

V Proofs and Extensions

This is a surprisingly easy theorem to prove. We will discuss several proofs here. The main analytical tool for proving equilibrium existence theorems are fixed point theorems. Fixed point theorems are of the form: For all sets X and functions $f : X \to X$ which satisfy some additional properties, there is an $x \in X$ such that $f(x) = x$. This is a neat trick for proving equations have solutions. For instance, the equation $f(x) = 0$ has a solution iff the function $g(x) = f(x) + x$ has a fixed point. Border (1985) is a good source on fixed point theorems.

The fundamental fixed point theorem is is due to Brouwer (1912):

Theorem 5 (Brouwer's Fixed-Point Theorem:). *If $K \subset \mathbf{R^n}$ is compact and convex, and if $f : K \to K$ is continuous, then f has a fixed point.*

Here is one proof of Theorem 1. This is modeled after the simplest proof I know of the existence of competitive equilibrium.

Proof of Theorem 1: In the game $\hat{G}$ strategy spaces are the closed unit simplices in $\mathbf{R}^{|\mathbf{S_i}|}$. Thus each S_i, and therefore S is compact and convex. For each player i, mixed strategy profile σ and pure strategy s_i, define

$$\Delta_i(s_i, \sigma) = U_i(s_i, \sigma_{-i}) - U_i(\sigma_i, \sigma_{-i}).$$

The function $\Delta_i(s_i, \sigma)$ measures how much better pure strategy s_i does than σ_i against σ_{-i}. Now define $\beta > 0$ such that

$$\beta^{-1} = 0.1 + \max_i \max_{s_{-i}} \Big(\max_{s_i} u_i(s_i, s_{-i}) - \min_{s_i} u_i(s_i, s_{-i}) \Big)$$

Thus β^{-1} is an upper bound on how big Δ_i can be. There is nothing significant about 0.1, except that it exceeds 0, so the upper bound is strict.

The trick is to define a function on $\hat{S}$ whose fixed points, if they exist, will be precisely the Nash equilibria. The function f will have one coordinate for every pure strategy in the game. Define f such that

$$f_{s_i}(\sigma) = \sigma_i(s_i)\big(1 + \beta\Delta_i(s_i, \sigma_{-i})\big). \tag{1}$$

The intuition behind this function is to think of a fictitious strategy adjustment process. For player i, if pure strategy s_i pays off more than average against σ_{-i}, the probability on s_i is raised. If s_i pays off less than average, it is lowered.

The trick is to define a function on $\hat{S}$ whose fixed points, if they exist, will be precisely the Nash equilibria. The function f will have one coordinate for every pure strategy in the game. Let $|S_i|$ denote the number of pure strategies belonging to player i. Define $f^{ni} : \hat{S} \to \mathbf{R}^{|\mathbf{S_i}|}$ such that

$$f^{ni}_{s_i}(\sigma) = \frac{\sigma_i(s_i)\big(1 + \beta\Delta_i(s_i, \sigma)\big) + \frac{1}{n}}{1 + \frac{1}{n}|S_i|}$$

This function is obviously continuous. We need to see that $f^{ni}(\sigma)$ is in fact in $\hat{S}_i$, which is a subset of $\mathbf{R}^{|\mathbf{S_i}|}$. To do this we need to check two things: First, that each $f^{ni}_{s_i}(\sigma) \geq 0$, and second, that $\sum_{s_i} f^{ni}_{s_i}(\sigma) = 1$. Each $f^{ni}_{s_i}(\sigma)$ is non-negative because we carefully constructed β so that the coefficient multiplying $\sigma_i(s_i)$ would be positive. In fact, you can see that each $f^{ni}_{s_i}(\sigma)$ is at least $1/(n + |S_i|) > 0$. To see the second claim, sum the numerators of the $f^{ni}_{s_i}(\sigma)$ to get $1 + \sum_{s_i} \sigma_i(s_i)\Delta(s_i, \sigma) + \big(|S_i|/n\big)$. The sum is the average of deviations from the mean, and so is 0. Check this by computing it from the definition of Δ_i.) Hence the ssum of the numerators equals the dominator, and so the sum $\sum_{s_i} f^{ni}_{s_i}(\sigma)$ is 1.

Now we define the map $f^n : \hat{S} \to \hat{S}$ such that

$$f^n(\sigma) = \big(f^{n1}(\sigma), \ldots, f^{nI}(\sigma)\big) .$$

The function f^n and the set $\hat{S}$ together satisfy all the conditions of Brouwer's Theorem, and so it has a fixed point σ^n. No fixed point σ^n will be a Nash equilibrium, because it must be a completely mixed strategy. But the sequence $\{\sigma^n\}$ is contained in a compact set, and so it has at least one subsequential limit $\sigma^* \in \hat{S}$. We will complete the proof by showing that any such limit σ^* is a Nash equilibrium.

We need again to show two things: First that for all i and $s_i \in S_i$, $\Delta(s_i, \sigma^*) \leq 0$. In other words, there is no pure strategy which does better against σ^*_{-i} that σ^*_i. Second, we must show that if $\Delta_i(s_i, \sigma^*) < 0$, then $\sigma^*_i(s_i) = 0$. Both of these facts follow from the fixed-point property of the σ^n, that

$$\Big(1 + \frac{1}{n}|S_i|\Big)\sigma^n_i(s_i) = \sigma^n_i(s_i)\big(1 + \beta\Delta_i(s_i, \sigma^n)\big) + \frac{1}{n} .$$

To show the first claim, notice that the fixed-point property implies that

$$\frac{1}{n}|S_i| > \beta\Delta(s_i, \sigma^n) .$$

As n becomes large, the left side converges to 0 and the right converges to $\Delta_i(s_i, \sigma^*)$, so taking limits, $0 \geq \Delta_i(s_i, \sigma^*)$.

To prove the second claim, compute from the fixed point property that

$$\sigma^n_i(s_i) = \frac{\frac{1}{n}}{\frac{1}{n}|S_i| - \beta\Delta_i(s_i, \sigma^n)} .$$

Suppose that $\Delta_i(s_i, \sigma^*) < 0$. Then the limit as $n \to \infty$ exists, and

$$\sigma_i^*(s_i) = \frac{0}{-\beta \Delta_i(s_i, \sigma^*)} = 0\,.$$

Thus σ^* is a Nash equilibrium. ☐

To check your understanding of the proof, show that the same method also works with the function

$$g_{s_i}(\sigma) = \frac{\max\{0, \sigma_i(s_i) + \Delta(s_i, \sigma_i)\}}{\sum_{s_i} \max\{0, \sigma_i(s_i) + \Delta_i(s_i, \sigma_i)\}}$$

We will use another approach to prove Theorem 2. Here is the idea: For every strategy profile s we can find the set of strategy profiles $F(s)$ with the property if $t \in F(s)$, then for all i, t_i is a best reply to s_{-i}. This generalizes the idea of reaction functions in the analysis of simple Cournot duopoly models. Suppose there was an s with the property that $s \in F(s)$. Then s_i is a best reply to s_{-i} for all i, and so s is a Nash equilibrium. This seems like a fixed point argument again, but instead of a function $f : S \to S$ we have a point-to-set map. Point-to-set maps are called correspondences.

Definition 6. *A correspondence $F : X \hookrightarrow Y$ is a function from X to 2^Y, the collection of all subsets of Y.*

We will work with a correspondence from S to itself, and again, look for fixed points. The fixed point theorem we will use is due to S. Kakutani, the father of the *New York Times* cultural affairs writer Michiko Kakutani (and a very distinguished mathematician). The theorem states conditions under which a correspondence $F : K \hookrightarrow K$ has a fixed point. As before, $K \neq \emptyset$ will have to be compact and convex. Also needed are continuity conditions for F.

Definition 7. *A correspondence $F : X \hookrightarrow Y$ has closed graph iff the set $\{(x, y) : x \in X, y \in F(x)\}$ is closed.*

In particular, this definition implies that for all $x \in X$, $F(x)$ is closed. To check your understanding of the definition, prove that if $f : X \to Y$ is a continuous function, it has a closed graph.

So far this is like Brouwer's theorem, which Kakutani's theorem directly generalizes. For functions, $f(x)$ is always a single point. For correspondences, $F(x)$ could be any set, and so we will need to control this, too.

Definition 8. *A point $k \in K$ is a fixed point of a correspondence $F : K \hookrightarrow K$ iff $k \in F(k)$.*

Theorem 6 (Kakutani's Fixed-Point Theorem). *If $K \subset \mathbf{R}^\mathbf{n}$ is compact and convex, if $F : K \hookrightarrow K$ has closed graph, and if for all $k \in K$, $F(k)$ is non-empty and convex, then F has a fixed point.*

Kakutani's theorem can be used to prove Theorem 2.

Proof of Theorem 2. Define $F_i : S \hookrightarrow S_i$ to be the correspondence such that

$$F_i(s) = \{t_i : \text{for all } r_i \in S_i,\, u_i(t_i, s_{-i}) \geq u_i(r_i, s_{-i})\}.$$

The set $F_i(s)$ is the set of best responses to s_{-i}. If u_i is quasi-concave in s_i for each s_{-i}, the set $F_i(s)$ is convex. Continuity of u_i implies that it is closed. Continuity and the compactness of S_i implies that it is not empty.

The fact that the correspondence $F_i(s)$ has closed graph is a consequence of the maximum principle, which you either have seen in 609 or will see in 610. Briefly, suppose $s^n \to s$ and $t_i^n \in F_i(s^n)$ and $t_i^n \to t_i$. We need to show that t_n is a maximizer of $u_i(\,\cdot\,, s_{-i})$. Suppose not. Then there is an $r \in S_i$ such that $u_i(r_i, s_{-i}) > u_i(t_i, s_{-i})$. Choose $\epsilon < \big(u_i(r_i, s_{-i}) - u_i(t_i, s_{-i})\big)/3$ and positive. By continuity, there is a $\delta > 0$ such that if $||s'_{-i} - s_{-i}|| < \delta$, $|u_i(r_i, s_{-i}) - u_i(r_i, s'_{-i})| < \epsilon$. There is also a $\delta' > 0$ such that if $||(t'_i, s'_{-i}) - (t_i, s_{-i})|| < \delta'$, $|u_i(t'_i, s'_{-i}) - u_i(t_i, s_{-i})| < \epsilon$. Since the sequence (t_i^n, s_{-i}^n) converges to (t_i, s_{-i}), it follows that for n large enough, $u_i(r_i, s_{-i}^n) - u_i(t_i^n, s_{-i}^n) > \epsilon/3$, and so $t_i^n \notin F_i(s^n)$, which is a contradiction.

The final step is to define $F(s) = \prod_{i \in I} F_i(s)$, so that $F : S \hookrightarrow S$. It is easy to check that F is convex- and nonempty-valued, and has closed graph. By hypothesis, S is the product of convex and compact sets, and hence itself convex and compact. Thus the conditions of Kakutani's theorem are met, and so there exists an $s^* \in S$ such that $s^* \in F(s^*)$. Thus $s_i^* \in F_i(s^*)$, so each s_i^* is a best reply to s_{-i}^*. That is, s^* is a Nash equilibrium. □

Suppose that the u_i are not quasi-concave, but merely continuous. Equilibrium may not exist. However, as we did in the case of finite games, we can define the mixed extension of G, wherein we allow for randomized strategies and use expected utility to define payoffs in the extended game. Players randomized strategies are independent, and expectations are linear in probability measures. So if we fix σ_{-i}, expected utility will be linear in σ_i. Linearity implies quasi-concavity, and so the set of best replies to any σ_{-i} is convex. For any set S_i, the set of probability distributions on S_i are convex. We don't know how to make sense of compactness and closed-graph in this setting, but there is a topology which makes this all work out. Consequently, the same proof technique can be used to prove the following theorem:

Theorem 7. *Let $G = (I, (S_i, u_i)_{i \in I}$ denote a game such that:*

1. *I is finite;*
2. *for all $i \in I$, S_i is a compact subset of a finite dimensional vector space;*
3. *for all $i \in I$, u_i is continuous.*

Then the mixed extension $\hat{G}$ has a Nash equilibrium.

For details, see Glicksberg (1952).

VI Bibliography and Extra Reading

Border, Kim C., *Fixed point theorems with applications to economics and game theory*, Cambridge: Cambridge University Press, 1985.

Glicksberg, I. L., "A Further Generalization of Kakutani's Fixed Point Theorem, with Application to Nash Equilibrium Points", *Proceedings of the American Mathematical Society* 3:1, 1952, pp. 170–174.

Harsanyi, J. C. and R. Selten, *A general theory of equilibrium selection in games*, Cambridge MA: MIT Press, 1988.

Schelling, Thomas, *The strategy of conflict*, London: Oxford University Press, 1963.