Nash Equilibrium

Nicola Giocoli

In 1950 the *Proceedings* of the National Academy of Sciences published a one-page paper by a twenty-one-year-old Ph.D. student in the Princeton mathematics department. The paper contained a new solution concept for noncooperative games with *n* players and no zero-sum constraint, called the equilibrium point. The concept was later dubbed Nash equilibrium after the name of its creator.

A *Nash equilibrium* (NE) is a collection of strategies by the *n* players such that no player can improve his outcome by changing only his own strategy. That is, a strategy profile $\hat{s} = (\hat{s}_1, \hat{s}_2, \ldots, \hat{s}_n)$ is an NE if, for every i and every strategy $s_i \in S_i$, we have: $u_i(\hat{s}) \geq u_i(\hat{s}_1, \ldots, \hat{s}_{i-1}, s_i, \hat{s}_{i+1}, \ldots, \hat{s}_n)$, where S_i is the set of strategies and u_i is the utility function of player i. The formal definition clarifies that the NE is a strategy profile in which the strategy played by each player is at least as good a reply as any other strategy available to him to the strategies played by the other players.

Hundreds of papers and books bear witness to the outstanding position of NE in economic theory. In 1994 Nash's contribution to economics was officially recognized by the awarding of the Nobel Prize. Apparently

Correspondence may be addressed to Nicola Giocoli, Department of Economics, University of Pisa, via Curtatone e Montanara 15, 56126 Pisa, Italy; e-mail: giocoli@mail.jus.unipi.it. I wish to thank Nicolò Bellanca, Mark Blaug, Marco Dardi, Luca Lambertini, Robert Leonard, Ivan Moscati, and Roy Weintraub for their useful comments and suggestions. I am of course fully responsible for any remaining mistakes. The financial support of MIUR PRIN 2002 "Mathematics in the History of Economics" is gratefully acknowledged.

History of Political Economy 36:4 © 2004 by Duke University Press.

the award was well deserved since, according to the standard opinion, NE is the embodiment of the idea that economic agents are rational, that they simultaneously act according to their incentives. Given that this idea is the driving force of economic theory, it follows that NE embodies no less than the most fundamental idea of economics.[1]

Yet, as is well known,[2] Nash was no economist at all: he was a mathematician and, as Roy Weintraub (1999, 212) put it, neither the Nobel Prize in economics nor the systematic "repetition of the phrase 'Nash equilibrium'" will ever make him an economist. Nash equilibrium came out as a purely mathematical answer to a purely mathematical question, that is, as an outgrowth of the fixed-point technique whose most immediate application, as suggested by its inventor, was to . . . poker! Thus, it would seem pointless to make Nash's work a topic for the history of economic thought and, in particular, to try to place NE in the context of postwar economic theory. However, the issue gains relevance as soon as it is recognized that the new concept was *not* an immediate hit in the economists' community. The popularity of game theory in general, and of NE in particular, is indeed a relatively recent event. In the 1950s and 1960s most neoclassical economists simply ignored that their discipline's central concept, rational equilibrium behavior, had finally found a precise, simple, and very general formulation. Even in the 1970s, game theory still remained a discipline for the specialists, and it was at least a decade away from making its official entry into the tables of contents of standard textbooks in economics.[3] A study of citation patterns based on the Social Science Citation Index and the Science Citation Index shows that in the eighteen years from 1966 to 1983 two key papers of modern game theory, Nash 1950b and von Neumann [1928] 1959, were cited, respectively, only sixty-three and fifty times (Oehler 1990, 106–7). Furthermore, the widely used textbook on game theory by Guillermo Owen ([1970] 1982) mentioned NE twice on page 127 and 143, and then never again for the remainder of more than three hundred pages! How could this happen?

Celebrating the fiftieth anniversary of NE, a leading game theorist has argued that the concept marked "one of the great watershed breakthroughs in the history of social science" because it modified the definition of economics itself, from the old image of a specialized science

---

1. See Aumann 1987, 43; and van Damme and Weibull 1995, 15.
2. See the best-selling biography by Sylvia Nasar (1998).
3. The true landmark was Kreps 1990a.

concerned with the production and allocation of material goods to the contemporary image of a science concerned with the analysis of incentives in all social institutions (Myerson 1999, 1067–68). He adds that "before Nash, price theory was the one general analytical methodology available to economics," while "the broader analytical perspective of noncooperative game theory has liberated practical economic analysis from this methodological restriction. . . . So Nash's formulation of noncooperative game theory should be viewed as one of the great turning points in the long evolution of economics and social science" (1080). The problem is that stressing the methodological revolution brought forth by NE can be hardly reconciled with the standard opinion that the concept was "just" the long-awaited "perfect" formulation of the concept of rational behavior *which had always been there* in economics, at least since the advent of the marginalist/neoclassical approach.

The goal of the following pages is to explain why NE had been neglected for many years by neoclassical economists, including the most mathematically (or even game-theoretically) oriented ones. I will argue that the main reason behind the neglect did not lie in the "revolutionary" character of the concept underlined by Myerson, but rather in its inconsistency with the research agenda of postwar mainstream economists. At the top of that agenda was the "how and why" of equilibrium, the explanation or justification of the way an economic system achieves an equilibrium position. This was a topic that the economists working in the 1950s had inherited from those active in the interwar years, who had raised the issue only to leave it unsolved. NE could offer no response either, for two reasons: first, the kind of rationality supporting it implied a modeling of the economic agent's epistemic capabilities that far exceeded what a postwar economist would find acceptable; and second, the static fixed-point view of equilibrium embodied by NE made it impossible to tackle the issue of how agents revised their plans and expectations—in other words, of the very topic that had been singled out since the 1930s as the crucial one in order to elucidate the how and why of equilibrium.

Acknowledging the gap separating NE from 1950s neoclassical economics is not just a way to answer the historical puzzle of the former's late reception. Actually, both reconstructing the theoretical context where the new notion should have found application and explaining why this did not happen also provide a significant contribution to *contemporary* economics, and one that only historians can give. It is not always recognized, in fact, that even today there does not exist unanimous consensus as to the interpretation of NE.[4] Regardless of the concept's formal elegance and analytical convenience, asking game theorists "To what question is NE the answer?" may still give rise to heterogeneous and, at best, only partial replies. Once more, how can this be possible, given that NE allegedly embodies nothing but the rigorous characterization of the notion of rational behavior *which has always been around* in the discipline?

From the historians' viewpoint, the enduring troubles in its interpretation show that NE, as well as the idea of economic man underlying it, really marked a dramatic discontinuity with respect to the orthodox notions of equilibrium and rationality. This provides a useful lesson as to what constitutes progress in economics. Nash equilibrium, born from the ingenious mind of a mathematician, harshly criticized during its first years, exploited by a few mathematical economists as a mere analytical tool for goals other than its natural ones, then forgotten for almost twenty years, was finally to be rediscovered, appreciated, and officially enthroned as *the* equilibrium concept and *the* characterization for rational behavior sought by generations of economists. Yet, the persistent difficulty in conflating NE with the theoretical edifice of which it should constitute the main foundation highlights the problems that economics has to face whenever its progress takes place by borrowing from other disciplines (usually, mathematics) notions and tools that might fail to harmonize with its core method, goals, and propositions.

1. Nash's Contributions to Noncooperative Game Theory

Nash wrote very few papers on game theory: seven in all, of which three were joint works, plus a twenty-seven-page Ph.D. dissertation.[5] The most famous one is the one-page paper "Equilibrium Points in n-Person Games," written in the autumn of 1949 and published the following year in the *Proceedings of the National Academy of Sciences* (Nash 1950b). This is the place where the NE was first formulated, and the work that gained the author his Nobel Prize. The three other papers written by Nash

---

[4] For recent, and not-so-recent, general assessments of this issue see, e.g., Johansen 1982; Aumann 1987; Kreps 1987, 1990b; Mas-Colell, Whinston, and Green 1995, 248–49; Jacobsen 1996; Mailath 1998; and Hargreaves Heap and Varoufakis 2004, chaps. 2–3.

[5] All the papers have been republished in Nash 1996. A photostatic reproduction of Nash's dissertation is in Kuhn and Nasar 2002.

alone were the 1950 *Econometrica* paper "The Bargaining Problem," which offered an axiomatic treatment of the classic exchange problem (Nash 1950a); the 1951 *Annals of Mathematics* paper "Non-cooperative Games," which reproduced with some changes Nash's Ph.D. thesis (Nash 1951); and finally the 1953 *Econometrica* paper "Two-Person Cooperative Games," which extended the analysis of the piece on bargaining to include the new setup of noncooperative game theory (Nash 1953). In what follows I will focus on the two papers on noncooperative games.[6]

1.1 The Equilibrium Point

In Nash 1950b the new solution concept, called the equilibrium point, was conceived of as a generalization of the existence result established by John von Neumann with the minimax theorem for two-person zero-sum games (2P ZSG). Both in a pioneering 1928 paper (von Neumann [1928] 1959) and in the 1944 book written with Oskar Morgenstern, the *Theory of Games and Economic Behavior*, von Neumann's goal had been to give a general mathematical characterization of rational behavior. He had concluded that to be rational in a strategic situation meant to play the *minimax strategy*: to choose the strategy that maximized the gain among the worst possible outcomes that might arise due to the rival's choice, or that minimized the maximum loss that the rival might cause. The powerful *minimax theorem* provided the formal support to this conclusion. However, the validity of von Neumann's characterization of rationality did not extend beyond the restricted class of 2P ZSG. Due to this limitation, von Neumann and Morgenstern tried to reduce all games with more than two players and/or without the zero-sum constraint to a 2P ZSG played between two coalitions of players. This forced them to limit the analysis to the *cooperative* setup, where coalitions could be formed and held together. Again, this fell short of fulfilling the goal of a truly general characterization of rationality.

That a renowned genius like von Neumann had failed to provide a satisfactory solution to a mathematical problem he himself had raised challenged the talent and wit of the young mathematicians working in the Princeton mathematics department in the late 1940s and early 1950s. The challenge was successfully met by the most talented of all, John Nash, who provided a new solution concept as well as a new general

---

[6] For general appraisals of Nash's contribution, see Leonard 1994 and Milnor 1998. See also the technical assessments by van Damme and Weibull (1995) and Binmore (1996).

approach for the reduction of all games to a *noncooperative* setup. In the last lines of his paper Nash (1950b, 49) explicitly remarked that the new concept was a generalization of the old minimax. Hence, from a strictly mathematical viewpoint, the result did represent an improvement: von Neumann's great achievement, the minimax theorem, had been reduced to a special case of Nash's equilibrium point.

The argument in the 1950 paper was crystal clear. Given an n-player game where each player has a finite set of pure strategies and can make recourse to the mixed strategies, any n-tuple of strategies $(s_1, s_2, \dots, s_n)$ is a point in the product space $S^n = S_1 \times S_2 \times \cdots \times S_n$, where each S_i is a finite-dimensional simplex. An n-tuple of strategies is said to *counter* (C) another "if the strategy of each player in the countering n-tuple yields the highest obtainable expectation for its player against the n-1 strategies of the other players in the countered n-tuple." Formally, $s^*C\hat{s}$ if, for every $i = 1, 2, \dots, n$, $s_i^* \in s^*$ is such that $u_i(s_i^*, \hat{s}_{-i}) \geq u_i(s_i, \hat{s}_{-i})$ for any $s_i \in S_i$, where S_i is the strategy set available to player i and s_{-i} is the (n-1)-tuple of strategies played by all other players. A self-countering n-tuple, that is, an n-tuple such that $u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*)$, $\forall i, s_i$, is called an *equilibrium point*.

To prove the existence of an equilibrium n-tuple, Nash observes that the correspondence of each n-tuple with its set of countering n-tuples gives a one-to-many mapping of the product space into itself, $S^n \rightarrow S^n$. The graph of this correspondence is closed due to the continuity of the payoff functions $u_i(\cdot)$. Moreover, the set of countering n-tuples of a given n-tuple is convex due to the linearity of the expected payoff function à la von Neumann and Morgenstern. Since we have a closed graph with convex images, it is possible to apply Kakutani fixed-point theorem to prove that the mapping has a fixed point, which is the desired equilibrium. Hence Nash proves that every n-player finite game played with mixed strategies admits an equilibrium point. Moreover, for the first time in game theory he establishes an explicit link between the notion of equilibrium and that of a fixed point.[7]

1.2 Noncooperative Games

The 1951 paper "Non-cooperative Games," published in the *Annals of Mathematics*, is the printed version, with few minor changes and one

---

[7] It is remarkable that the word *equilibrium* is seldom used in the *Theory of Games*. Thus it is only with Nash that the term really entered the language of game theory and it is only in recent times that it has become synonymous with "solution to a game."

major omission, of the Ph.D. thesis that Nash completed in the spring of 1950. From the very first lines of the work Nash underlines the distance separating his own approach from that of von Neumann and Morgenstern. The latter are said to have developed a theory of n-person cooperative games; Nash's theory, instead, is based upon the absence of coalitions: each player is assumed to behave independently, without any collaboration or communication with the other players (Nash 1951, 286).

The basic ingredient of the new theory is the notion of the equilibrium point, which is once more presented as a generalization of von Neumann's solution for 2P ZSG. Yet, the offered definition is not anymore in terms of countering, as it coincides with the modern one: an n-tuple $s = (s_1, \ldots, s_n)$ is an equilibrium point if, for every i, we have: $p_i(s) = \max_{r_i} p_i(s, r_i)$, where $p_i(\cdot)$ is player i's payoff and $p_i(s, r_i) = p_i(s_1, \ldots, s_{i-1}, r_i, s_{i+1}, \ldots, s_n)$. In words: "An equilibrium point is an n-tuple s such that each player's mixed strategy maximizes his payoff if the strategies of the others are held fixed. Thus each player's strategy is optimal against those of the others" (287).

The major theoretical innovation of the 1951 paper is the concept of *solution*.[8] A game is said to be *solvable* if its set S of equilibrium points satisfies the following *interchangeability condition*: $(t, r_i) \in S$ and $s \in S$ imply $(s, r_i) \in S, \forall i$, that is, if it is possible to replace the strategy played by player i in the equilibrium n-tuple s with the one played in the equilibrium n-tuple t and still obtain an equilibrium n-tuple. The solution of a solvable game is the set S of its equilibrium points (290). A noncooperative game does not always have a solution, but when it has one, the solution is unique. Nash also defines *strong solvability*: a game is strongly solvable if it has a solution S such that, for all i, $s \in S$ and $p_i(s, r_i) = p_i(s)$ imply $(s, r_i) \in S$, that is, if every unilateral deviation from an equilibrium n-tuple causing no change in the payoff still determines an equilibrium n-tuple (290). Finally, Nash defines the *value* of a game. Let $v_i^+ = \max_{s \in S}[p_i(s)]$ and $v_i^- = \min_{s \in S}[p_i(s)]$ be the upper and lower values of the game. If the two coincide, then v is the value of

---

[8] Another remarkable feature is the proof used to demonstrate the existence of an equilibrium point for every n-player game. Nash abandons the Kakutani theorem employed in the previous paper and adopts the Brouwer fixed-point theorem. What is rather peculiar in this replacement is Nash's comment that the new proof constitutes "a considerable improvement over the earlier version" (1951, 288). Since Kakutani theorem is a generalization of Brouwer, the comment is quite mysterious. For a possible explanation, see Giocoli 2003c, 304–5. For the history of how the various fixed-point theorems entered neoclassical economics, see Giocoli 2003b.

the game (291). Of course, whenever there is just one equilibrium point, the value always exists.

These definitions show that what Nash meant by the term *solution* as applied to a noncooperative game, was not an equilibrium point of the game, but rather a *set* of equilibrium points such that the equilibrium strategies of the players were interchangeable. This is because what he was looking for was, exactly like von Neumann and Morgenstern in the *Theory of Games*,[9] an *objective* configuration of the payoffs that might represent the game's value, that is, the amount each agent might expect to get by playing the game. However, while von Neumann and Morgenstern's solution, the minimax, did characterize objectively a strategic situation since, by choosing the minimax strategy, each player freed himself from the need to foresee the rival's moves, thereby making the play really independent of any *subjective* feature such as beliefs, expectations, etc., Nash's solution was *not* completely objective, as it still depended on the latter features by requiring every player to correctly anticipate the rival's equilibrium move (see next section).

The 1951 paper ended with a section highlighting some of the directions for future research. The most obvious one was the analysis of all n-person games "for which the accepted ethics of fair play imply noncooperative playing" (294), as, for instance, in poker. A less obvious direction was the analysis of cooperative games. In Nash's definition, a *cooperative game* is a strategic situation where players can and will collaborate as they did in the *Theory of Games*. Thus, in such a game the players can communicate and form coalitions that will be enforced by an umpire (295). However, in the cooperative setup à la von Neumann and Morgenstern the payoffs were assumed to be freely transferable, and all compensations and agreements were assumed to take place outside the game, during the so-called preplay negotiations. Nash considered these assumptions unnecessarily restrictive and proposed a new approach where "any desired transferability can be put into the game itself instead of assuming it possible in the extra-game collaboration" (295). This passage heralded what later came to be known as *Nash's program* in game theory. The new approach required in fact that the pre-play negotiations be modeled in such a way that the negotiation steps themselves became moves in a larger noncooperative game with an infinite number of pure strategies. By doing so, the problem of investigating a cooperative game turned into that of obtaining a suitable noncooperative

---

[9] See, e.g., von Neumann and Morgenstern [1944] 1953, 31, 77, 160.

model for the negotiation phase. The larger game could then be analyzed according to the theory of the equilibrium point, extended to the infinite case. Whenever a value for such a game existed, it also was the value for the original cooperative game (295). This was of course a major breakthrough with respect to von Neumann and Morgenstern's approach, as it allowed the whole cooperative theory to be subsumed under the more general noncooperative setup. Nash himself will give an excellent example of the new approach in the 1953 *Econometrica* paper on cooperative games—the last of his great contributions to game theory.

1.3 The Omitted Section

Nash's Ph.D. thesis was not entirely reproduced in the *Annals of Mathematics*. The dissertation ended in fact with a section on the interpretation of equilibrium points that disappeared from the published version. Yet, this "ghost" section is a fundamental element for both the reconstruction of Nash's own views and the interpretation of NE. Its title was "Motivation and Interpretation" and its goal was that of showing how equilibrium points and solutions (in the above-specified sense) could be connected with observable phenomena. These two pages therefore represent the only available evidence of Nash's effort to give a *positive* reading of his results.[10]

Nash suggested two possible interpretations of equilibrium points. The first is the *"mass-action" interpretation*, which requires that many rounds of the same game be played. In this dynamic setup there is no need to postulate that players have full knowledge of the game structure or any special reasoning ability. Agents are instead assumed to be able to exploit their experiences in previous rounds of the game in order to accumulate information about the relative merits of the available strategies. In short, the players are not modeled as perfectly rational, fully informed "demi-gods,"[11] but rather as *learning* human beings.

In order to show that such a characterization of the participants in the game may indeed lead to an equilibrium point, Nash made three assumptions. The first is that there is a statistical population of players for each possible position of the game. The second is that the average round of the game involves *n* players selected at random from the *n* populations. The

---

[10] The omitted section has been reprinted in Nash 1996, 32–33, and Kuhn and Nasar 2002, 78–81.

[11] I borrow this expression from Morgenstern [1935] 1976, 173.

third is that each pure strategy is employed by the "average member" of a population according to a stable average frequency. Since the game is noncooperative, there is no collaboration among the agents playing in the different positions. It follows that the probability that a given n-tuple of pure strategies be played in a round of the game is the product of the chances that each of the n pure strategies be played in a random playing of the game.

The tricky issue is that of modeling the learning process, that is, of formalizing the effect of past experience on the players' behavior. Nash elegantly solved it by assuming that the agents' accumulation of empirical evidence on the pure strategies at their disposal can be captured by postulating that those who occupy the position i of the game learn the numbers $\mathrm{p}_{\mathrm{i}\alpha}(\mathrm{s}) = \mathrm{p}_{\mathrm{i}}(\mathrm{s}, \pi_{\mathrm{i}\alpha})$, that is, the expected payoffs belonging to someone in the ith position who plays the pure strategy $\pi_{\mathrm{i}\alpha}$. The key observation is that if the agents learn these numbers, they will only play optimal pure strategies: the pure strategies $\hat{\pi}_{\mathrm{i}\alpha}$ such that $\mathrm{p}_{\mathrm{i}\alpha}(\mathrm{s}) = \max_\beta \mathrm{p}_{\mathrm{i}\beta}(\mathrm{s})$ (Nash 1996, 32). It follows that in the mixed strategy $\mathrm{s}_{\mathrm{i}} = \sum_\alpha \mathrm{c}_{\mathrm{i}\alpha}\pi_{\mathrm{i}\alpha}$ which results from the average play of player i, a positive probability weight $\mathrm{c}_{\mathrm{i}\alpha}$ will be attached only to optimal pure strategies $\hat{\pi}_{\mathrm{i}\alpha}$. Thus, the pure strategy $\hat{\pi}_{\mathrm{i}\alpha}$ is used in s_{i} only if $\mathrm{p}_{\mathrm{i}\alpha}(\mathrm{s}) = \max_\beta \mathrm{p}_{\mathrm{i}\beta}(\mathrm{s})$. But the latter is nothing but the condition for $\mathrm{s} = (\mathrm{s}_1, \mathrm{s}_2, \ldots, \mathrm{s}_{\mathrm{n}})$ to be an equilibrium point. Therefore, "the assumptions we made in this 'mass-action' interpretation lead to the conclusion that the mixed strategies representing the average behavior in each of the populations form an equilibrium point" (33). In other words, the average player can *learn the equilibrium behavior*; he can converge to a kind of playing where only the equilibrium mixed strategies are employed on the average. Nash added that if all the assumptions hold, the argument does not even need to be based upon the existence of a large population.

As to the positive value of this interpretation, Nash observed that "there are situations in economics or international politics in which, effectively, a group of interests are involved in a non-cooperative game without being aware of it; the non-awareness helping to make the situation truly non-cooperative" (33). Yet, he acknowledged that the best we can expect to observe in reality is just an approximation of equilibrium, since the spread and utilization of information on past rounds of the game as well as the stability of the average frequencies will seldom show the kind of perfection required by the assumptions. With the benefit of hindsight, it is quite straightforward to recognize in Nash's mass-action interpretation a forerunner of the modern *evolutionary* argument for NE.[12]

The second interpretation suggested by Nash should sound even more familiar to modern game theorists. In this interpretation, which applies to games played only once, the notion of solution (that is, of equilibrium points with interchangeable strategies) plays a major role. The starting point is "the question: what would be a 'rational' prediction of the behavior to be expected of rational playing of the game in question?" (33). Nash argued that by using the principles that "a rational prediction should be unique, that the players should be able to deduce and make use of it, and that such knowledge on the part of each player of what to expect the others to do should not lead him to act out of conformity with the prediction, one is led to the concept of a solution defined before" (33).

The positive tone of this interpretation was given by the following remark. Let $S_1, S_2, \ldots, S_n$ be the sets of equilibrium strategies of a solvable game. Then, according to Nash, the rational *prediction* of the game should be that a rational player occupying the position *i* would on the average play the mixed strategy s_i in S_i. This in fact is what would turn out "if an experiment were carried out" (33). Hence, the second interpretation was said to positively capture the outcome of a would-be empirical framework reproducing the strategic situation.[13]

Nash concluded by observing that the second interpretation was "quite strongly a rationalistic and idealizing" one. In order for the players to be able to deduce the prediction for themselves, the assumption must be made in fact that they fully know the structure of the game. While even this interpretation was buried in the ghost section, it did not take long for other game theorists to propose it, so that it quickly became *the* explanation of NE.

2. Interpreting Nash Equilibrium

The standard definition of neoclassical equilibrium as a set of rational, simultaneous, and mutually compatible (or *consistent*) plans had been

---

[12] See Mailath 1998, 1354–63. The pioneering work on evolutionary games is that by Maynard Smith (1982). Van Damme and Weibull (1995, 34–37), Milnor (1998, 1331), and, above all, the Nobel Committee (see Royal Swedish Academy of Sciences 1995, 3) all emphasize this feature of Nash's contribution. Another early "evolutionary" argument may be found in Luce and Raiffa 1957, 105.

[13] Working at the RAND Corporation Nash will carry out one such experiment, albeit mainly focusing on cooperative games; see Kalish, Milnor, Nash, and Nering 1954. See also Luce and Raiffa 1957, 259–69.

developed in the 1930s by, among others, Hayek, Lindahl, and Hicks.[14] On account of it, one of the most important issues for neoclassical economics soon became that of understanding the constraints that the rationality postulate placed upon the elaboration and revision of the agents' plans, that is, upon the process through which the economic system reached an equilibrium. Let us call this issue the *learning problem* of neoclassical economics. As I have argued at length elsewhere,[15] most major economists of the interwar period—including the above-mentioned inventors of the new notion of equilibrium—recognized that finding an answer to the learning problem was the key to attaching a *positive* content to neoclassical theory.

Today it is straightforward to extend the general definition of neoclassical equilibrium to situations of strategic interdependence. The problem is that its game-theoretic counterpart, the NE, is usually given a purely formal interpretation as a static fixed point that leaves no room for the issue of how and why the players learn to play it in the first place. In order to tackle the learning problem, the players should be modeled as only boundedly rational. This is what Nash actually did with his "mass-action" interpretation. However, such an ingenious solution could only be conceived by someone who was completely free of the preconceptions about the agents' rationality typical of neoclassical economists—in short, by someone who, like John Nash, was no economist at all. Hence, with the disappearance of the ghost section from Nash 1951 also went the possibility of showing the economists' community that a rigorous explanation of the how and why of equilibrium could indeed be achieved. What will be argued in the following pages is precisely that the failure of NE to influence neoclassical economics for almost thirty years was mainly due to the circumstance that of the two interpretations that Nash gave of his concept, the one that might have raised the interest of those postwar economists who were struggling with the learning problem remained unknown, while the one that was to become popular among game theorists, although nicely fitting the *modern* (i.e., end-of-twentieth-century) image of hyperrational, perfectly foreseeing agents, had no chance of being appreciated by most 1950s economists, as it negated the very essence of their privileged research topic.

[14] See Hayek 1937, Lindahl 1939, and Hicks 1939.

[15] See Giocoli 2003c, chap. 3.

2.1 The Limits of the Standard Interpretation

Two major assumptions provide the foundations for noncooperative game theory.[16] The first is the hypothesis of *maximization*, the fundamental postulate of rationality. The second is the hypothesis of *consistency*: that the agents' expectations of the other agents' behavior are correct. The latter assumption implies that the overall pattern of individual optimizing choices constitutes an NE: from a decision-theoretic point of view, in fact, the definition of NE requires that the players be *correct* in their expectations about the strategies played by the other agents. Yet, the consistency hypothesis is *not* a condition for NE, but rather one of its defining features, the other being that stemming from the maximization hypothesis, namely, the absence of any gain from unilateral deviations. A major challenge for noncooperative game theory is to provide a compelling justification for these two hypotheses, and thus also for NE. In short, why should people play NE?

As far as the maximization hypothesis is concerned, the role of the assumption is to economize over the relevant characteristics of individual behavior, so that the analysis can focus on the institutional features of the strategic situation. If agents are "perfect" maximizers, then any change in their behavior can be interpreted in terms of the modifications of the institutional system of incentives and disincentives, and not in terms of shifts or imperfections in the agents' psychology or tastes. The maximization postulate was accepted by von Neumann and Morgenstern as well as by Nash, but none of them believed that it could suffice to characterize strategic behavior. Therefore a crucial role in noncooperative game theory is played by the second assumption, the hypothesis of consistency. Yet the most common explanation of NE emphasizes only the maximization side, while downplaying the role of the consistency side.

Given the noncooperative premise, player *i* selects his own strategy s_i to maximize his payoff given the strategies independently chosen by the other players $s_{-i} = (s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_n)$. There is a *best reply* that a rational agent can play against the other agents' strategies, but the player also expects his (rational) rivals to do the same, that is, to try in turn to play their best replies. This creates the well-known chain of conjectures "I think that he thinks that I think . . . " that makes the individual choice of the best strategy indeterminate. The chain can only be broken if there is a strategy $\hat{s}_i$ for each agent *i* satisfying the defining

---

[16] Here I follow Mailath 1998, 1347–48.

condition of NE. Thus, the existence of an NE n-tuple $\hat{s}$ is necessary to avoid the standstill caused by the infinite regress of conjectures. This explanation highlights the fixed-point nature of NE, since it turns out that the equilibrium n-tuple is made of strategies, each of which is a best reply to the best replies to itself.

The interpretative problems of NE begin when we try to go beyond the previous explanation and figure out the how and why of NE. In other words, to what question is NE indeed the answer? According to David Kreps, the majority of modern game theorists share the following position. Assume that, by some unspecified means, an agreement has been reached by the players on how each of them will play the game. Then NE is a *necessary* condition for the agreement to be *self-enforcing*: only if the agreement calls for each player to play his own NE strategy is the agreement itself *stable*, that is, no single player has the incentive to deviate unilaterally from it. The condition, however, is far from sufficient to establish a socially stable situation, as we see immediately if we consider multi-player deviations. Moreover, this story does not explain how the agreement comes about in the first place. Finally, it does not clarify what happens if no agreement is reached.[17]

The standard explanation needs therefore to be integrated, if not fully replaced. Kreps deals in particular with the issue of how an agreement can be reached. One possibility is that of assuming explicit pre-play negotiations. The result is that, although we can guarantee neither that the players will reach an agreement nor what specific agreement will be reached, we can still be sure that the range of possible self-enforcing agreements, arrived at via pre-play negotiations, is contained within the set of NE. The problem is to identify a mechanism for pre-play negotiations. The nature of the agreement, and so of NE, will then depend on the specific mechanism. In case no explicit pre-play negotiation is possible, there are other possible explanations of how an agreement may be reached. Sometimes a player may know what the others will do (because, say, there exists a unanimously adopted theory, known by everybody, of how to play the game), or at least be extremely confident in the validity of his belief upon their actions. If this holds for all players, a sort of implicit agreement arises. This kind of explanation falls under the headline of *focal points*: if there exists an obvious way to play the game, every player will know what the others are doing.[18]

---

[17] See Kreps 1987, 584; and 1990b, 28–36.

[18] See Kreps 1990b, 143–44; and Schelling 1956, [1960] 1980.

Yet, according to Kreps, none of these and similar stories are ever fully convincing. The difficulties with the standard explanation and with its most immediate extensions help explain the success of an alternative approach to game theory that gives much more weight to the consistency assumption: the *Bayesian approach*.

2.2 Bayesian Game Theory

Traditionally, Bayesian decision theory had been perceived as appropriate only for tackling exogenous, rather than strategic, uncertainty. However, starting in the late 1970s, a rich literature has explored noncooperative game theory from an explicit decision-theoretic Bayesian viewpoint. Central to this literature is the assumption that players assign subjective probabilities to *all* uncertainty, including the actions and beliefs of other players. Games are then analyzed in terms of the players' rationality and of their *epistemic* states, namely, what they know or believe about the game and about each other's rationality, actions, knowledge, and beliefs. The various solution concepts are distinguished according to the epistemic content (basically, common knowledge assumptions) of the players' subjective probabilities. Today, most leading game theorists believe that game theory is nothing but the extension of decision theory to the case of two or more decision makers. Since the standard characterization of rationality in decision-making problems is the Bayesian one, it follows that the Bayesian approach is essential even for game theory: "there is no way that a book on game theory can be written if Bayesian rationality cannot be assumed" (Binmore 1992, 119).

More specifically, Bayesian game theorists set out to investigate where the assumption of the players' rationality could lead in a noncooperative game, and what further hypotheses about the players' epistemic endowments were necessary to get to NE. As I said before, Bayesian players are required to quantify via subjective probability distributions *all* the uncertainty they face—both that depending on the states of nature and that depending on the other players' choices and beliefs. Moreover, these distributions must be common knowledge among the players: they must be known, known to be known, and so on, ad infinitum.

These tight requirements upon the players' epistemic attributes provide the desired compelling justification for the hypothesis of consistency: players are correct in their conjectures about the other agents' strategy choices *because* every agent's beliefs are common knowledge.

Furthermore, when it is assumed that all players know the other players' conjectures about each other's behavior, a proper epistemic foundation for NE obtains. In other words, if the players' payoffs, their rationality, and their beliefs are all common knowledge, then the profile of the players' beliefs constitutes an NE. This central result, demonstrated by Tan and Werlang (1988), gives an epistemic characterization of NE as an *equilibrium of subjective beliefs*, rather than of strategies.[19]

With such a purely subjective account of the equilibrium, Bayesian game theory has reached in a sense the opposite end of the spectrum with respect to von Neumann and Morgenstern's objective characterization of a game solution (see above, sec. 1.2). Yet, the fact remains that the Bayesian approach has seemingly succeeded in explaining the how and why of NE. Or, has it?

2.3 Strategic Uncertainty and Epistemic Endowments

Bayesian game theory can offer us an important lesson with respect to the issue of the interpretation of NE, namely, that in order to "justify" the choice of the NE strategies it is necessary to model the players as endowed with exceptional epistemic requirements. The players must be rational, must know the game structure and must know the other players' ways of thinking. Moreover, all this information must be common knowledge. It follows that if we share the standard view that rationality as understood in game theory amounts to playing the NE strategy,[20] the characterization of rational play must include these exceptional epistemic abilities. The point is that this "justification" of NE betrays the spirit of the "founding fathers" of modern game theory—von Neumann, Morgenstern, and Nash. Hence, historically speaking, the idea that strategic rationality coincides with Bayesian rationality is highly questionable.

All the three founding fathers believed that strategic uncertainty was totally different from stochastic uncertainty, and thus could not be tackled

---

[19] For further results, see Brandenburger and Dekel 1989 and Aumann and Brandenburger 1995. In particular, the latter offers a slightly less demanding result, namely, that if the players share a common prior, if the payoffs and the players' rationality are mutual knowledge, and if the players' beliefs are common knowledge, then the beliefs' profile constitutes an NE in mixed strategies (1163). Hence, a crucial common-knowledge assumption is still placed upon the players' conjectures.

[20] See Binmore and Dasgupta 1986, 8; and Aumann 1987, 43.

with the tools of probability theory, not even of the Bayesian kind. On the one side, von Neumann and Morgenstern unambiguously refused to characterize strategic rationality in Bayesian terms. According to them, the other players' choices could never be considered as chance events and reduced to subjective probability distributions because every player should always take into account the fact that such choices were the outcome of his opponents' will, which might be as rational as his own. Moreover, von Neumann and Morgenstern believed that standard decision problems should be considered as special cases of strategic analysis (i.e., as one-player games) and not the other way around.[21] On the other side, Nash too believed that it was impossible to model strategic uncertainty through subjective probabilities, and thus he never pursued the Bayesian path. The shared refusal to model the players as having subjective probabilities over the other agents' choices had serious analytical consequences in that it entailed that the founding fathers had to deal with the problem of strategic uncertainty by making recourse either to an objective solution concept (von Neumann and Morgenstern's minimax) or to the possibility of an implicit or explicit agreement over the strategies to be played (Nash's equilibrium point).

Yet, it has to be stressed that von Neumann and Morgenstern were far more explicit in refusing to found their analysis upon any particular epistemic endowment of the players. Indeed, they clearly stated in several places in the *Theory of Games* that what they were looking for was an *objective* characterization of rational behavior, that is, either one that could be authoritatively prescribed and explained to the players by an "umpire," or one that every player could understand by merely looking at the payoff structure of the game. From this point of view Nash did mark a break, because his new solution concept failed to fully comply with von Neumann and Morgenstern's goal (see above, sec. 1.2), so that it paved the way, although largely unintentionally, to the modern epistemic approach.

3. Interpreting Nash's Interpretations of NE

In this last section I argue that the different "justifications" offered by the founding fathers to the minimax and the NE can explain, first, the

---

[21] For the first point, see, e.g., von Neumann and Morgenstern [1944] 1953, 99. For the second, see the same work, page 86.

extent to which the Bayesian approach may be said to have provided a satisfactory interpretation of the latter concept, and, second, the reason why Nash's brilliant idea did not meet with immediate success in the economists' community.

3.1 From Solution to Equilibrium

Faced with the central problem of strategic uncertainty, that is, the problem of the infinite regress of conjectures, Nash refused to employ the subjective probabilities and reduce the game-theoretic setup to a special case of Bayesian decision theory. He dealt instead with the infinite regress by making recourse to an equilibrium concept. In other words, he solved strategic uncertainty by imposing an *equilibrium condition*.

The idea itself of a solution concept as a way out from the infinite regress problem is quite relevant for us. Modeling strategic, as well as parametric, uncertainty through subjective probabilities is a way to answer the central question of neoclassical theory, namely, the "how and why" of equilibrium. Understanding how agents formulate and revise their probabilistic beliefs constitutes a possible, albeit partial,[22] approach to the learning problem. This is not so in the realm of game theory, where strategic uncertainty is circumvented through equilibrium conditions. Thus, the notion of equilibrium plays a different role in game theory: instead of being the central category around which the analysis of the learning process can be organized and developed, it is a theoretical tool that avoids having to be concerned with the learning problem in the first place. In a sense, therefore, game theory uses equilibrium to beg the fundamental question of neoclassical economic theory.

This interpretation covers both von Neumann and Morgenstern's as well as Nash's approaches. For both versions of game theory the fundamental question is what it means to be rational in a game. For both, strategic rationality is the output and not the input (i.e., an assumption) of the analysis. For both, there is no need to explicitly model the agents' expectations. For both, the answer to the previous question is found by collapsing the characterization of rational behavior upon the respective solution concepts. This makes it impossible for both approaches to theorize

---

[22] The Bayesian model seems to be at present the only available "learning" model consistent with the cornerstones of neoclassical economics. Yet it is questionable whether it can be considered real learning what is just a consistent process of adaptation and updating of the agent's beliefs to new information.

over the learning process, that is, over the how and why of equilibrium. It follows that in both approaches the equilibrium conditions are not considered in relation to the learning process, that is, as emerging from a process of revision of plans and conjectures. The crucial difference between the founding fathers is that, while this omission causes no problem in the *Theory of Games*, it lies at the heart of the still problematic interpretation of NE. The break came when game theory passed from the normative idea of the game *solution* (namely, if players behave so and so, the interdependence is broken and a certain result is warranted) to the positive idea of the game *equilibrium*. Nash represented the watershed between the two notions: he aimed at *solving* all games, but to do so he devised an equilibrium notion. The NE did settle one of the two crucial questions of equilibrium, namely, that of its definition, but exacerbated the other one, namely, that of its interpretation—the how and why issue. The latter's importance was perceived by Nash himself, who tried to tackle it in the ghost section of his dissertation.

Interpreting the NE has become an ever more pressing issue as the notion of "solution" has become more and more in the modern economists' and game theorists' jargon synonymous with that of "equilibrium." Two attitudes have emerged. One is that of most neoclassical economists who have simply begged the issue by focusing on the defining features of NE as a static fixed point. The other is that of the Bayesians game theorists who have solved the issue by modeling the agents' epistemic endowments. As I said before, the latter attitude betrays the spirit of early modern game theory; yet, what matters most is that *both* attitudes betray the spirit and main theoretical agenda of interwar neoclassical economics. As a matter of fact, by adopting the equilibrium concept suggested by Nash, most economists have been forced to cut all ties with the typical themes of the 1930s and early 1940s, such as learning, imperfect foresight, disequilibrium dynamics. From Nash onward, all the emphasis has been either on the defining characteristics of equilibrium or on the players' epistemic capabilities that may "justify" it, while the crucial topic of the adjustment process leading to the equilibrium has instead been largely neglected.

3.2 Why the NE Was Neglected

The previous subsection helps us clarify why the NE did not meet with immediate success among economists. Borrowing Ken Binmore's terminology (1996, x-xiii), the standard "justification" of NE can be deemed a *negative* one: nothing other than an NE can be the solution of a game, or, as it is sometimes said, an authoritative game theory book cannot possibly recommend a strategy profile as a solution to a game unless it is an NE. Thus, if a noncooperative game has a solution, it must be among the NE of the game. However, when thinking of how to "justify" the choice of one strategy instead of another, it is natural for a player to look for a *positive* justification. This is what happens in the case of von Neumann's 2P ZSG: a rational player has a positive reason for choosing the maximin strategy because it is the strategy that guarantees him a certain payoff. Thus, if asked to explain why he played that particular strategy, or if called to prescribe a rational rule of conduct, an agent may always convincingly argue in favor of the maximin choice. This does not apply in the case of the standard justification for NE.

Yet, Nash did not neglect the problem of providing a positive justification for his equilibrium concept. Of the two interpretations of equilibrium points offered in the ghost section of his Ph.D. thesis, only the second one was fairly in line with the now standard "negative" argument for NE, while the first one was of a truly positive kind. As we know, the "mass-action" interpretation was based upon an iterative adjustment process, in which boundedly rational players observed the strategies played by their opponents randomly drawn from a uniformly distributed population of players, and gradually learned to adjust their own strategies to get higher payoffs. Nash suggested that the learning process would eventually converge to an NE, if it converged to anything at all, and remarked that in this interpretation it was unnecessary to assume that players had a full knowledge of the game structure or the ability to go through any complex reasoning process. Thus, the logical steps of the mass-action interpretation entailed a dynamics that could account for the how and why of equilibrium.

It is remarkable that in the ghost section, and only there, Nash embraced a descriptive viewpoint. His two interpretations aimed in fact at showing "how equilibrium points and solutions can be connected with observable phenomena" (Nash 1996, 32). Consequently, he seemed to believe that the mass-action interpretation could describe the actual behavior of real economic agents in strategic situations. It may therefore be argued that, at least as far as the mass-action interpretation is concerned, Nash's *own* game theory did not cut all ties with the central issue of many interwar neoclassical economists, that is, with the learning problem as the key to providing economic theory with an empirical content.[23]

The absence of the two interpretations from the published version of Nash's thesis produced the curious situation of a would-be breakthrough economic concept, the NE, appearing in a mathematical journal deprived of the very features that would have made it *really* and *immediately* interesting to the economists' community, that is, an original solution to the crucial question of the learning process. What remained in the 1951 paper was the new solution concept and the formal proof of its existence. The latter, being based on a fixed-point theorem, had clearly a "negative" flavor that clashed with the research agenda of most postwar neoclassical economists where the theoretical puzzles of the interwar period still featured prominently.

Moreover, independently of Nash's thesis and even before its full epistemic requirements were revealed by modern Bayesians, the second interpretation of NE—the one consistent with the fixed-point view—was quickly put forward as *the* interpretation of the new concept.[24] The point is that this interpretation was also immediately perceived, as Nash himself had put it, as "quite strongly a rationalistic and idealizing" one, on account of its heroic assumptions about the agents' intellectual capabilities—something that the most acute commentators did not fail to notice, and forcefully reject, as early as the 1950s.[25] Add to this the fact that the two papers that Nash published in *Econometrica*—a journal that was more commonly found in economics departments than the *Annals of Mathematics*—were for obvious reasons those dealing with the bargaining problem, where fewer references, if any, to the new solution concept could be found. Sum all these negative elements together, and take into account the generally low level of mathematical literacy of postwar economists, and you will come really close to understanding why NE, the alleged embodiment of that very notion of rationality that had always been there in economics, went almost unnoticed for nearly three decades *as a characterization of rational behavior*.[26]

---

[23] In this sense the flourishing subdiscipline of evolutionary game theory (see above, footnote 12) seems to bring game theory back to tackling the fundamental issues of interwar economics. This, it should be noted, notwithstanding the explicit anti-neoclassical stance of most evolutionary game theorists.

[24] See, e.g., Luce and Raiffa 1957, 63.

[25] See Shubik 1952, 146; and Hurwicz 1953, 403.

[26] Not of course as a useful mathematical tool, as testified by the use made of it by Arrow and Debreu (1954, 273–79).

3.3 What History Can Do for Us: Rebutting an Ahistorical Claim

An opposite claim has been made by Ken Binmore, who has argued that if Nash's positive interpretation of NE had been known, the success of the concept would have been *undermined.* The mass-action interpretation was in fact too similar to the adjustment process of the Cournot duopoly model, and the latter had been harshly criticized in the economic literature of the period precisely because of its reliance on the firms' myopic behavior.[27] Thus, according to Binmore (1996, xii), what made NE palatable to postwar economists was that the way it was presented in the literature "freed them of the need they had previously perceived to tie down the dynamics of the relevant equilibrating process before being able to talk about the equilibrium to which it will converge in the long run." In other words, the NE allowed the tricky question of the learning process to be cleared up and forgotten, thereby overcoming in a single step what had looked like an unbridgeable gulf to past generations of economists.

My previous exposition shows that Binmore's remark is correct in all but its conclusion. The trouble lies in the ahistorical attitude with which he depicts the economists' point of view. What the post-Nash history of economics and game theory shows is that the lack of a positive justification for NE did not motivate the latter's success but rather its *delayed acceptance*. Far from welcoming any effort to wipe out the learning problem from economic theory, the average economist of the period could not appreciate an equilibrium concept lacking a positive interpretation and based upon the formal logic of the fixed-point argument. A clear example is given by oligopoly theory, the field closer to game-theoretic themes, where the trend from the late 1930s on had been that of tackling the how and why question by privileging empirical observations and field studies.[28]

But the best confirmation of my argument is given by the way Martin Shubik reinterpreted the Cournot equilibrium of the duopoly model with the goal, to use Robert Leonard's phrase, of "giving Nash a history."[29] The motive for Shubik's new reading of Cournot has to be found in his change of attitude toward the NE: after having repeatedly denied the relevance of the new solution concept, Shubik elected to make it the

---

[27] The classic critique to the Cournot model was that by Fellner (1949, chap. 2).

[28] See Mason 1939, Katona 1946, Rothschild [1947] 1953, and Fellner 1949. I have reconstructed part of this story in Giocoli forthcoming.

[29] See Leonard 1994, 507–9.

cornerstone of his 1959 book, *Strategy and Market Structure*. Yet, being one of the few economists who really knew game theory, he realized that the NE was so alien to standard neoclassical analysis that in order to found upon it his new theory of oligopolistic competition he had better anchor the concept to the works of some noble forefather.

In the eyes of the only economists who really worked side by side with the great Princeton mathematicians, the messages to communicate to the economists' community were, first, that the equilibrium point was all that mattered when analyzing the cases of "competition among the few"; second, that the conventional approach focusing on the adjustment to equilibrium—with the all-too-known difficulties implicit in any effort to explain the how and why of equilibrium—could be confidently abandoned; and third and most important, that the newly suggested approach had a tradition in economic analysis. As a consequence, what until then had been the standard (and, by the way, basically correct) view of Cournot's solution as being based upon a dynamic process of myopic adjustments[30] was transformed in a single stroke. According to Shubik (1959, 63), Cournot's solution could be given an entirely *static* interpretation: "We can interpret [Cournot's] equilibrium point as being such that if the duopolists were at such a point neither would be motivated to change his rate of production. *This does not explain how or why the individuals move to this point.* It merely specifies the equilibrium property once it has been attained" (emphasis added).

So, far from being potentially damaged by its similarity with one of the classic solutions to duopolistic competition, NE was deliberately "given a history" by linking it with Cournot's name, with the specific intent of making it palatable to economists.[31] What Shubik failed to consider, however, was that, although the tastes for equilibrium analysis of a very small and geographically concentrated group of mathematical economists[32] were indeed changing in the period, the majority of the economists' community—especially, I repeat, those in the crucial field of oligopoly theory—still adhered to the traditional pre-World War II view whereby the how and why of equilibrium was really the fundamental issue. With the benefit of hindsight, it may be argued that Shubik would have done better to recover Nash's mass-action interpretation and

---

[30] See Magnan de Bornier 1992, 625–31; and Giocoli 2003a, 183–87.

[31] That NE had no relation to the standard interpretation of Cournot's solution in terms of myopic adjustment along the reaction curves had already been noted by Hurwicz (1953, 402).

[32] Namely, those working at the Cowles Commission.

propose it as a formal argument proving that Cournot, the "true" Cournot of myopic adjustments, was after all "almost" right . . .

Shubik's misperception of the demands arising from the economists' community may thus have contributed to the failure of NE to stimulate the analytical fantasy of the profession. As Leonard (1994, 508), put it: "Shubik's writings capture perfectly the gradual shift away from the implicit acceptance of Fellner's critique, in which the implausibility of the reaction curve dynamics is emphasized at the expense of the equilibrium, towards a static interpretation in which the dynamics are suppressed, and the equilibrium point preserved." Shubik's reinterpretation of Cournot marks therefore the true landmark for the application of game-theoretic, static equilibrium theory to economic analysis. That this happened by "kidnapping" Cournot from the pantheon of the forerunners of neoclassical dynamic analysis to enthrone him as the patron saint of modern noncooperative theory is another useful indication of what constitutes progress in economic theory.

4. Conclusion

The main solution concept of noncooperative game theory—the NE—is the keystone of the whole theoretical edifice of contemporary neoclassical economics.[33] Thus, it is no exaggeration to claim that *today* NE embodies the discipline's most fundamental notion, namely, that of rational strategic behavior. Yet, despite its importance and widespread use, the concept still lacks a compelling interpretation; it still fails to explain how and why the participants in a noncooperative game should come to play their equilibrium strategies. From an historical viewpoint, the main puzzle of the first decades of modern game theory is why the new discipline in general, and NE in particular, were almost completely neglected by mainstream economics. What I have tried to show in this essay is that this puzzle may really be explained only by taking into account the interpretation issue, in both its contemporary and Nash's own version.

Yet, there is a more general lesson that this story may tell us. Consider neoclassical economics immediately after World War II. The discipline was locked in a stalemate that dated back to the interwar years and that originated in the economists' failure to transform their discipline into a truly empirical science on account of their inability to satisfy the

---

[33] This section draws on Giocoli 2003c, 333–36.

necessary prerequisite for that transformation, namely, the modeling of the agents' learning process and of the dynamic equilibrating path. My claim is that the most remarkable feature of Nash's game theory was precisely that it contained a solution to such a stalemate. Indeed, both his new solution concept (the NE) and one of its interpretations (the converging adjustment process implemented by boundedly rational individuals) were immediately translatable in terms of the economists' learning problem.

Of course, we can only surmise what the development of postwar neoclassical economics would have been if it had managed to benefit from Nash's ingenious idea. If we consider that these were also the years when, to mention just one name, Herbert Simon was starting his long journey along the route of bounded rationality, it may even be suggested that the discipline might have eventually transformed itself into a real analysis of the social interaction of individuals which would combine a strong analytical foundation with a keen empirical interest.

Unfortunately, we know that Nash cut the two pages of the ghost section from the 1951 *Annals of Mathematics* paper. A key role in this editorial choice might have been played by the different evaluation of what is scientifically relevant for a mathematician rather than for an economist—indeed, the analytical gist of the paper was given by the section on solvability and the new Brouwer-based existence proof.[34] The lesson to be drawn for the history of economics is that the episode epitomizes the moral that "there are various incommensurabilities associated with the movements of ideas across disciplinary boundaries, and nowhere is this more apparent than in the movement of ideas from the community of mathematicians to the community of economists" (Weintraub 1999, 211).

References

Aumann, R. J. 1987. What Is Game Theory Trying to Accomplish? In *Frontiers of Economics*, edited by K. Arrow and S. Honkapohja, 28–76. Oxford: Basil Blackwell.

Aumann, R. J., and A. Brandenburger. 1995. Epistemic Conditions for Nash Equilibrium. *Econometrica* 63:1161–80.

Binmore, K. 1992. *Fun and Games: A Text on Game Theory*. Lexington, Mass.: D.C. Heath.

———. 1996. Introduction to Nash 1996.

---

[34] See section 1.2 above, and Kuhn and Nasar 2002, 48.

Binmore, K., and P. Dasgupta. 1986. Game Theory: A Survey. In *Economic Organisations as Games*, edited by K. Binmore and P. Dasgupta, 1–45. Oxford: Basil Blackwell.

Brandenburger, A., and E. Dekel. 1989. The Role of Common Knowledge Assumptions in Game Theory. In *The Economics of Missing Markets, Information, and Games*, edited by F. Hahn, 46–61. Oxford: Oxford University Press.

Fellner, W. 1949. *Competition among the Few*. New York: Knopf.

Giocoli, N. 2003a. “Conjecturizing” Cournot: The Conjectural Variations Approach to Duopoly Theory. *HOPE* 35.2:175–204.

———. 2003b. Fixing the Point: The Contribution of Early Game Theory to the Tool Box of Modern Economics. *Journal of Economic Methodology* 10.1:1–39.

———. 2003c. *Modeling Rational Agents: From Interwar Economics to Early Modern Game Theory*. Cheltenham, U.K.: Elgar.

———. Forthcoming. The Escape from Conjectural Variations: The Consistency Condition in Duopoly Theory from Bowley to Fellner. *Cambridge Journal of Economics*.

Hargreaves Heap, S., and Y. Varoufakis. 2004. *Game Theory: A Critical Text*. 2nd ed. London: Routledge.

Hayek, F. A. 1937. Economics and Knowledge. *Economica* 4:33–54.

Hicks, J. R. 1939. *Value and Capital: An Inquiry into Some Fundamental Principles of Economic Theory*. Oxford: Clarendon Press.

Hurwicz, L. 1953. What Has Happened to the Theory of Games. *American Economic Review* 43.2:398–405.

Jacobsen, H. J. 1996. On the Foundations of Nash Equilibrium. *Economics and Philosophy* 12:67–88.

Johansen, L. 1982. On the Status of the Nash Type of Noncooperative Equilibrium in Economic Theory. *Scandinavian Journal of Economics* 84.3:421–41.

Kalish, G. K., J. W. Milnor, J. F. Nash Jr., and E. D. Nering. 1954. Some Experimental n-Person Games. In *Decision Processes*, edited by R. M. Thrall, C. H. Coombs, and R. L. Davis, 301–27. New York: John Wiley.

Katona, G. 1946. Psychological Analysis of Business Decisions and Expectations. *American Economic Review* 36:44–62.

Kreps, D. 1987. Nash Equilibrium. In *The New Palgrave: A Dictionary of Economics*, edited by J. Eatwell, M. Milgate, and P. Newman, 584–88. London: Macmillan.

———. 1990a. *A Course in Microeconomic Theory*. New York: Harvester-Wheatsheaf.

———. 1990b. *Game Theory and Economic Modelling*. Oxford: Clarendon Press.

Kuhn, H. W., and S. Nasar, eds. 2002. *The Essential John Nash*. Princeton: Princeton University Press.

Leonard, R. J. 1994. Reading Cournot, Reading Nash: The Creation and Stabilisation of the Nash Equilibrium. *Economic Journal* 104:492–511.

Lindahl, E. 1939. The Dynamic Approach to Economic Theory. In *Studies in the Theory of Money and Capital*, 21–136. London: George Allen and Unwin.

Luce, R. D., and H. Raiffa. 1957. *Games and Decisions*. New York: John Wiley.

Magnan de Bornier, J. 1992. The "Cournot-Bertrand Debate": A Historical Perspective. *HOPE* 24.3:623–56.

Mailath, G. J. 1998. Do People Play Nash Equilibrium? Lessons from Evolutionary Game Theory. *Journal of Economic Literature* 36:1347–74.

Mas-Colell, A., M. D. Whinston, and J. R. Green. 1995. *Microeconomic Theory*. Oxford: Oxford University Press.

Mason, E. S. 1939. Price and Production Policies of Large-Scale Enterprises. *American Economic Review* 29.1 (supplement): 61–74.

Maynard Smith, J. 1982. *Evolution and the Theory of Games*. Cambridge: Cambridge University Press.

Milnor, J. 1998. John Nash and *A Beautiful Mind*. *Notices of the American Mathematical Society* 45:1329–32.

Morgenstern, O. [1935] 1976. Perfect Foresight and Economic Equilibrium. In *Selected Economic Writings of Oskar Morgenstern*, edited by A. Schotter, 169–83. New York: New York University Press.

Myerson, R. B. 1999. Nash Equilibrium and the History of Economic Theory. *Journal of Economic Literature* 37:1067–82.

Nasar, S. 1998. *A Beautiful Mind*. New York: Simon and Schuster.

Nash, J. F., Jr. 1950a. The Bargaining Problem. *Econometrica* 18:155–62.

———. 1950b. Equilibrium Points in n-Person Games. *Proceedings of the National Academy of Sciences* 36:48–49.

———. 1951. Non-cooperative Games. *Annals of Mathematics* 54:286–95.

———. 1953. Two-Person Cooperative Games. *Econometrica* 21:128–40.

———. 1996. *Essays in Game Theory*. Cheltenham, U.K.: Edward Elgar.

Oehler, K. 1990. Speaking Axiomatically: Citation Patterns to Early Articles in General Equilibrium Theory. *HOPE* 22.1:101–12.

Owen, G. [1970] 1982. *Game Theory*. Orlando: Academic Press.

Rothschild, K. W. [1947] 1953. Price Theory and Oligopoly. In *Readings in Price Theory*, edited by G. J. Stigler and K. E. Boulding, 440–64. London: George Allen and Unwin.

Royal Swedish Academy of Sciences. 1995. The Nobel Memorial Prize in Economics 1994. *Scandinavian Journal of Economics* 97:11–7.

Schelling, T. C. 1956. An Essay on Bargaining. *American Economic Review* 46:281–306.

———. [1960] 1980. *The Strategy of Conflict*. Cambridge: Harvard University Press.

Shubik, M. 1952. Information, Theories of Competition and the Theory of Games. *Journal of Political Economy* 60:145–50.

———. 1959. *Strategy and Market Structure: Competition, Oligopoly, and the Theory of Games*. New York: John Wiley.

Tan, T., and S. Werlang. 1988. The Bayesian Foundations of Solution Concepts of Games. *Journal of Economic Theory* 45:370–91.

van Damme, E., and J. W. Weibull. 1995. Equilibrium in Strategic Interaction: The Contributions of John C. Harsanyi, John F. Nash, and Reinhard Selten. *Scandinavian Journal of Economics* 97:15–40.

von Neumann, J. [1928] 1959. On the Theory of Games of Strategy. In vol. 4 of *Contributions to the Theory of Games*, edited by D. Luce and A. W. Tucker, 13–42. Princeton: Princeton University Press.

von Neumann, J., and O. Morgenstern. [1944] 1953. *Theory of Games and Economic Behavior*. Princeton: Princeton University Press.

Weintraub, E. R. 1999. Review of *A Beautiful Mind*, by Sylvia Nasar. *Journal of the History of Economic Thought* 21:209–11.