SSE Working Paper Series in Economics
No. 2019:1

An elementary axiomatization of the Nash equilibrium concept

Mark Voorneveld [a]

[a] Stockholm School of Economics
Department of Economics

An elementary axiomatization of the Nash equilibrium concept

Mark Voorneveld*

Department of Economics, Stockholm School of Economics, Box 6501, 113 83 Stockholm, Sweden, mark.voorneveld@hhs.se

January 17, 2019

Abstract

For strategic games, the Nash equilibrium concept is axiomatized using three properties: *(i)* if the difference between two games is 'strategically irrelevant', then their solutions are the same; *(ii)* if a player has a strategy with a constant payoff, this player need not settle for less in any solution of the game; *(iii)* if all players agree that a certain strategy profile is optimal, then this strategy profile is a solution of the game.

Keywords: Nash equilibrium; axiomatization; solution concept
JEL codes: C72

SSE Working Paper Series in Economics 2019:1

*I thank Peter Sudhölter, William Thomson, and Dries Vermeulen for helpful discussions. Financial support by the Wallander-Hedelius Foundation under grant P2016-0072:1 is gratefully acknowledged.

1 Introduction

The Nash equilibrium concept is applicable in a wide range of applications, no matter if strategy spaces are large, small, embellished with additional topological or measure-theoretic structure. This note's main result (Theorem 3.2) therefore provides an easy axiomatization of the Nash equilibrium concept for many classes of games.

Let me briefly discuss its three axioms. It is common — precise references to related literature are in Section 2, where the axioms are discussed in detail — to treat the two games

	L	R
T	$3,2$	$0,1$
B	$1,3$	$2,4$

	L	R
T	$3+a, 2+c$	$0+b, 1+c$
B	$1+a, 3+d$	$2+b, 4+d$

with real parameters a, b, c, and d as strategically equivalent. For instance, since the same vector (a, b) of payoffs is added to both rows of the row player's payoff matrix, there is nothing she can do about their difference: whatever distinguishes her preferences in the two games is outside the row player's control. A similar argument applies to the column player. So in passing from the left game to the right, things like the players' pure or mixed best-response correspondences, weak and strict dominance relations, and rationalizable strategies are all unaffected. The first axiom, *independence of identical consequences*, formalizes this standard notion of strategic invariance: their solutions are the same.

Consider now the game

	L	R
T	$2, 2+c$	$-2, 1+c$
B	$0, 3+d$	$0, 4+d$

obtained by setting a to -1 and b to -2. This gives the row player a risk-free option: playing the second row always gives payoff zero, no matter what the column player does. The second axiom, *risk-free back-up*, says that she need not accept less than this payoff in any solution of the game.

And if we set $(a, b, c, d) = (2, 0, 3, -3)$, we obtain the game

	L	R
T	$5,5$	$0,4$
B	$3,0$	$2,1$

Since both players strive to maximize their utility, they unanimously agree that (T, L) is the best strategy profile. The final axiom, *unanimity*, says that strategy profiles that are optimal for all players simultaneously are indeed solutions of the game.

Using a strategic invariance assumption — *independence of identical consequences*— to relate the solutions of different games departs from the existing literature on axiomatizations of the Nash equilibrium concept where this role is played by a consistency axiom. The notion of *consistency* for solution concepts of noncooperative games was introduced by Peleg and Tijs (1996) and Peleg et al. (1996) and requires the following. Assume a strategy profile x is picked out by the solution concept for a given game. Some players agree to play according to this strategy profile and leave the game. Then the remaining players don't have an incentive to deviate from their recommendations either: the restriction of x to the remaining players is a solution of the 'reduced game' obtained by plugging the leaving agents' assigned strategies into the utility functions. The reader is referred to Thomson (2001, Sec. 12) for an extensive overview of the axiomatic literature in noncooperative games.

Definitions, axioms, and a discussion of these axioms are in Section 2. Section 3 contains the main axiomatization, together with a more general discussion of relations between the different axioms and their logical independence. Variants — of the domain and/or the axioms — and concluding remarks are in Section 4.

2 Definitions and axioms

A (strategic or normal-form) *game* $G = \langle N, (X_i)_{i \in N}, (u_i)_{i \in N} \rangle$ has a nonempty, finite set N of players; each player $i \in N$ has a nonempty set X_i of strategies and a utility function $u_i : \times_{j \in N} X_j \to \mathbb{R}$ representing this player's preferences over strategy profiles. Function $\exp u_i / (1 + \exp u_i)$, with values between zero and one, is an order-preserving transformation of the utility function u_i. Since an order-preserving transformation represents the same preferences, we assume without loss of generality that all utility functions are bounded.

Let Γ be a set of games. A *solution concept* on Γ is a function 'sol' that assigns to each game $G = \langle N, (X_i)_{i \in N}, (u_i)_{i \in N} \rangle$ in Γ a subset $\mathrm{sol}(G)$ of its strategy space $X = \times_{j \in N} X_j$. We characterize the solution concept mapping each game G to its possibly empty set of Nash equilibria:

$$\mathrm{Nash}(G) = \{x \in X : \text{ for each } i \in N \text{ and } y_i \in X_i,\ u_i(x) \geq u_i(y_i, x_{-i})\}.$$

Here, as usual, $x_{-i} \in X_{-i} = \times_{j \in N \setminus \{i\}} X_j$ is the profile of strategies of i's fellow players and (y_i, x_{-i}) is the strategy profile obtained from x after i unilaterally deviates to $y_i \in X_i$.

We now introduce and briefly discuss our main axioms.

> *Independence of identical consequences:* If two games $G = \langle N, (X_i)_{i \in N}, (u_i)_{i \in N} \rangle$ and $H = \langle N, (X_i)_{i \in N}, (v_i)_{i \in N} \rangle$ in Γ differ only in their utility functions, but consequences of each player's behavior on her utility are the same, i.e.,
>
> for all $i \in N$, all $x_i, y_i \in X_i$, and all $x_{-i} \in X_{-i}$:
> $$u_i(x_i, x_{-i}) - u_i(y_i, x_{-i}) = v_i(x_i, x_{-i}) - v_i(y_i, x_{-i}), \quad (1)$$
>
> then their solutions are the same: $\mathrm{sol}(G) = \mathrm{sol}(H)$.

This is a standard invariance assumption (see below) to express that the difference between two games is 'strategically irrelevant': for each player i, whatever distinguishes her utility functions in the two games is *outside her control* and objects — whenever well-defined — like her (pure or mixed) best-response correspondence, (weak and strict) dominance relations, and rationalizable strategies are the same in the two games. In a noncooperative game, player i can't affect the strategy profile x_{-i} of the remaining players. And given that strategy profile, (1) says that i cannot affect the difference between u_i and v_i, because it is independent of what strategy she chooses:

$$\text{for all strategies } x_i \text{ and } y_i \text{ of player } i: \qquad u_i(x_i, x_{-i}) - v_i(x_i, x_{-i}) = u_i(y_i, x_{-i}) - v_i(y_i, x_{-i}). \quad (2)$$

That is, the difference $u_i(\cdot, x_{-i}) - v_i(\cdot, x_{-i})$ is constant, no matter what i does! We follow the convention from Facchini et al. (1997, p. 195) and Voorneveld et al. (1999, Sec. 2) and say that games G and H differ by a *dummy game*, a game where the utility of each player is unaffected by her own choice.

The strategic equivalence of games satisfying (1) is fundamental to at least four strands of literature. Firstly, in the context of equilibrium selection and refinement, Wu and Jiang (1962, p. 1310) call such games 'isomorphic'. Harsanyi and Selten (1988, pp. 77-80) discuss the strategic equivalence of such games in the chapter 'Consequences of Desirable Properties'. Peleg et al. (1996, p. 90) refer to such games as 'equivalent'. Secondly, in the literature on learning and evolution, this strategic invariance condition is used to describe games with similar behavior under dynamic processes; see for instance Blume (1993, Def. 6.2) and Weibull (1995, p. 19). Thirdly, it is used to characterize games that are strategically indistinguishable from games where all players have the same interests in the literature on potential games; see Monderer and Shapley (1996, Sec. 2), Morris and Ui (2004, Def. 3), or Voorneveld (2010, p. 405). Fourthly, the decision-theoretic irrelevance of such transformations of the utility functions was already observed in the early axiomatic literature on decision theory under uncertainty; see Milnor (1954, axiom 7).

The next axiom says that if a player has a strategy with a constant utility of — say — c, then she doesn't need to settle for less, because she can assure getting c all on her own:

> *Risk-free back-up:* If, in a game $G = \langle N, (X_i)_{i \in N}, (u_i)_{i \in N} \rangle$ in Γ, a player i has a strategy y_i with a constant utility, i.e.,
>
> $$\text{there is a } c \in \mathbb{R} \text{ with } u_i(y_i, x_{-i}) = c \text{ for all } x_{-i} \in X_{-i},$$
>
> then each solution of the game gives i at least this payoff: $u_i(x) \geq c$ for all $x \in \mathrm{sol}(G)$.

This constant is a lower bound on the player's minimax/individually rational payoff:

$$c \leq \inf_{x_{-i} \in X_{-i}} \sup_{x_i \in X_i} u_i(x_i, x_{-i}). \tag{3}$$

Indeed, for each $x_{-i} \in X_{-i}$,

$$c = u_i(y_i, x_{-i}) \leq \sup_{x_i \in X_i} u_i(x_i, x_{-i}),$$

making c a lower bound on the function $x_{-i} \mapsto \sup_{x_i} u_i(x_i, x_{-i})$. So its *greatest* lower bound, the infimum over the strategies of i's fellow players, satisfies (3). Therefore, *risk-free back-up* is implied by the more demanding individual rationality assumption familiar from, for instance, zero-sum games and various folk theorems for repeated games where the minimax value provides lower bounds on achievable outcomes:

> *Individual rationality:* Players receive at least their minimax value of the game:
>
> $$\text{for all } G = \langle N, (X_i)_{i \in N}, (u_i)_{i \in N} \rangle, i \in N, x \in \mathrm{sol}(G): \qquad u_i(x) \geq \inf_{y_{-i} \in X_{-i}} \sup_{y_i \in X_i} u_i(y_i, y_{-i}).$$

But *individual rationality* is more demanding: it says something about each game in Γ, whereas *risk-free back-up* imposes restrictions only on games where a strategy gives a player a constant utility.

The final axiom simply says that if all players agree that a strategy profile is optimal, then it is a solution of the game:

> *Unanimity:* For each game $G = \langle N, (X_i)_{i \in N}, (u_i)_{i \in N} \rangle$ in Γ, if there is a strategy profile $x \in X$ with $u_i(x) = \max_{y \in X} u_i(y)$ for all players i, then $x \in \mathrm{sol}(G)$.

3 Axiomatization and relations between the axioms

A set of games Γ is *closed under strategic equivalence* if for each game G in Γ and each game H that differs from G by a dummy game, as in (1) or (2), the game H also lies in Γ. Some examples:

Example 3.1. Under the traditional convention that the set of players and their strategies all belong to some well-defined set of potential players/strategies, the following domains are closed under strategic equivalence: the set of all games, the set of all finite games (each X_i is finite: only pure, no mixed strategies), and the set of all games with at least one Nash equilibrium. But also — in contrast with axiomatizations based on *consistency*— the subdomains of the previous three obtained by keeping the player set and the set of strategies of each player fixed, varying only their utility functions. Moreover, since the potential games of Monderer and Shapley (1996) and supermodular games (see, e.g. Zhou, 1994, Sec. 3) are defined in terms of restrictions on the payoff differences

$$u_i(x_i, x_{-i}) - u_i(y_i, x_{-i})$$

and these are the same, by (1), for games that differ by a dummy game, also these two common classes of games give rise to domains that are closed under strategic equivalence. $\lhd$

Section 4 treats some other domains. The axiomatization of the Nash equilibrium concept is:

Theorem 3.2. *Let Γ be a set of games that is closed under strategic equivalence. The Nash equilibrium concept is the unique solution concept on Γ that satisfies independence of identical consequences, risk-free back-up, and unanimity.*

The Nash equilibrium concept satisfies the three axioms. For *independence of identical consequences*, games satisfying (1) have the same best-response correspondences and consequently the same Nash equilibria. And Nash equilibria give each player at least her minimax value:

$$\text{sol}(G) \subseteq \text{Nash}(G) \text{ for all } G \in \Gamma \implies \text{sol satisfies } \textit{individual rationality}.$$

When discussing the axioms, we already argued

$$\text{sol satisfies } \textit{individual rationality} \implies \text{sol satisfies } \textit{risk-free back-up}.$$

Also, strategy profiles maximizing each player's utility function are Nash equilibria, so

$$\text{sol}(G) \supseteq \text{Nash}(G) \text{ for all } G \in \Gamma \implies \text{sol satisfies } \textit{unanimity}.$$

So it remains to show that the Nash equilibrium concept is the *only* solution concept satisfying *independence of identical consequences*, *risk-free back-up*, and *unanimity*. That follows from:

Lemma 3.3. *Let Γ be a set of games that is closed under strategic equivalence. If a solution concept* sol *on Γ satisfies:*

(a) independence of identical consequences and risk-free back-up, then it assigns to each game a subset of its Nash equilibria: $\text{sol}(G) \subseteq \text{Nash}(G)$ *for all $G \in \Gamma$.*

(b) independence of identical consequences and unanimity, then it assigns to each game a set containing all Nash equilibria: $\text{Nash}(G) \subseteq \text{sol}(G)$ *for all $G \in \Gamma$.*

Proof. (a) Let solution concept sol satisfy *independence of identical consequences* and *risk-free back-up*. Let $G = \langle N, (X_i)_{i\in N}, (u_i)_{i\in N}\rangle \in \Gamma$ be a game and let $y \in \mathrm{sol}(G)$. To show: $y \in \mathrm{Nash}(G)$.

So let $j \in N$ be a player and $z_j \in X_j$ a potential deviation; why is $u_j(y) \geq u_j(z_j, y_{-j})$? Define game $H = \langle N, (X_i)_{i\in N}, (v_i)_{i\in N}\rangle$ with $v_i = u_i$ if $i \neq j$ and $v_j(x) = u_j(x) - u_j(z_j, x_{-j})$ for all $x \in X$. The functions $(v_i)_{i\in N}$ are bounded and G and H differ by a dummy game; since Γ is closed under strategic equivalence, game H is a well-defined game in Γ. By *independence of identical consequences*, $\mathrm{sol}(G) = \mathrm{sol}(H)$. So $y \in \mathrm{sol}(H)$. In game H, player j's strategy z_j gives her constant utility 0, so by *risk-free back-up*, $v_j(y) = u_j(y) - u_j(z_j, y_{-j}) \geq 0$.

(b) Let solution concept sol satisfy *independence of identical consequences* and *unanimity*. Let $G = \langle N, (X_i)_{i\in N}, (u_i)_{i\in N}\rangle \in \Gamma$ be a game and let $y \in \mathrm{Nash}(G)$. To show: $y \in \mathrm{sol}(G)$.

Let game $H = \langle N, (X_i)_{i\in N}, (v_i)_{i\in N}\rangle$ have utility functions defined for each $i \in N$ and $x \in X$ by

$$v_i(x) = u_i(x) - \sup_{z_i \in X_i} u_i(z_i, x_{-i}).$$

Firstly, each v_i is bounded, since each u_i is bounded. And G and H differ by a dummy game. Since Γ is closed under strategic equivalence, H is a well-defined game in Γ. Secondly, since y is a Nash equilibrium of G we have $v_i(y) = 0$ for all players i, whereas $v_i(x) \leq 0$ for all other strategy profiles x. So y maximizes each player's payoff in H. By *unanimity*, $y \in \mathrm{sol}(H)$. Thirdly, $\mathrm{sol}(G) = \mathrm{sol}(H)$ by *independence of identical consequences* and therefore $y \in \mathrm{sol}(G)$, as desired. □

The logical independence of the axioms is addressed in the next two examples.

Example 3.4 (Logical independence). The solution concept that assigns to each finite game:

1. the empty set satisfies all axioms in Theorem 3.2 except *unanimity*;
2. its entire set of strategy profiles satisfies all axioms except *risk-free back-up*;
3. its strong Nash equilibria satisfies all axioms except *independence of identical consequences*.

Recall that a strategy profile is a strong Nash equilibrium (Aumann, 1959) if no subset of players can deviate, keeping the strategies of the remaining players fixed, and make all its members better off. For instance, according to *independence of identical consequences*, the two games

	L	R
T	1, 1	3, 0
B	0, 0	2, 2

	L	R
T	4, 4	3, 3
B	3, 0	2, 2

must have the same solution, since the right game is obtained from the left by adding dummy game

	L	R
T	3, 3	0, 3
B	3, 0	0, 0

But the right game has strong Nash equilibrium (T, L) and the left game has none: *independence of identical consequences* is violated. ◁

The three solution concepts in the previous example establish logical independence of the axioms on several other nontrivial domains as well. But if the domain Γ of games is trivial — for instance if it contains only one-player games — the axioms are dependent:

Example 3.5 (One-player games). If Γ is a set of one-player games, then only *risk-free back-up* and *unanimity* are needed to axiomatize the Nash equilibrium concept. If G is a one-player game, each strategy of its single player is trivially risk-free. So *risk-free back-up* implies that each element of $\text{sol}(G)$ maximizes the player's utility function and is consequently a Nash equilibrium of G. Conversely, each Nash equilibrium maximizes the single player's utility function and must therefore belong to $\text{sol}(G)$ by *unanimity*. ◁

4 Variants and concluding remarks

Other domains; negative results. The axiomatization in Theorem 3.2 takes the domain to be closed under strategic equivalence. If this assumption is removed, the result need no longer be true. I give two examples. As a first extreme example, look at domain Γ_1 containing only one game: the mixed extension of

	L	R
T	$0,1$	$3,0$
B	$1,2$	$2,3$

Then *independence of identical consequences* has no bite since there are no distinct games in our class; *risk-free back-up* imposes no restrictions since none of the players has a mixed strategy with a constant utility; *unanimity* imposes no constraints since there is no strategy profile that simultaneously maximizes both players' utilities. In summary, in this singleton class of games, the three axioms impose no constraints on what the solution of this game should be.

Next, take the domain Γ_2 of mixed extensions of finite strategic games: the classical scenario of Nash (1950). Lemma 3.3(a) (inclusion '$\text{sol} \subseteq \text{Nash}$') goes through without change. But *independence of identical consequences* and *unanimity* are of little help to establish the opposite inclusion. In Lemma 3.3(b), these axioms are used to construct an auxiliary game H where a certain Nash equilibrium is payoff-maximizing for all players simultaneously. In the small domain Γ_2, such a construction is impossible in games without pure-strategy Nash equilibria. Why? Well, if a mixed-strategy profile gives each player her maximal utility, then all pure-strategy profiles in its support (those played with positive probability) must do so as well: they are pure-strategy Nash equilibria. In fact, there are solution concepts other than the Nash equilibrium concept on Γ_2, like

$$\text{sol}(G) = \begin{cases} \text{Nash}(G) & \text{if } G \text{ has pure-strategy Nash equilibria,} \\ \emptyset & \text{otherwise,} \end{cases}$$

that satisfy all three axioms in Theorem 3.2. So on the domain Γ_2 we would need substantially different axioms to characterize Nash equilibria, something outside the scope of the present paper. On the one hand, this is a potentially interesting topic for future research; on the other hand, Norde et al. (1996) already provide such a characterization using *consistency*.

Other domains; positive results. There are other common domains where the axiomatization *does* remain valid: instead of assuming that the domain is closed under strategic equivalence — if it contains a game G, then it contains *all* games H that differ from G by a dummy game — the proof of Lemma 3.3 goes through if for each game G in Γ the two *specific* auxiliary games H constructed in that proof lie in Γ as well. This is the case for the class Γ^* of games $\langle N, (X_i)_{i\in N}, (u_i)_{i\in N}\rangle$ where each X_i is a nonempty, convex, compact subset of some $\mathbb{R}^n$ with its usual topology, and each u_i is continuous and — given the strategy profile of i's opponents — quasiconcave in i's own strategy. This is the class of games for which Kakutani's fixed-point theorem easily assures the existence of Nash equilibria.

Indeed, let G be a game in this class. Look at the auxiliary game H in Lemma 3.3(b) with utility functions

$$v_i(x) = u_i(x) - \sup_{z_i \in X_i} u_i(z_i, x_{-i}).$$

By Berge's maximum theorem, the function $x \mapsto \sup_{z_i \in X_i} u_i(z_i, x_{-i})$ is continuous. Hence v_i, the difference between two continuous functions, is continuous as well. And given the strategy profile x_{-i} of i's fellow players, it simply subtracts a constant from u_i, so v_i remains quasiconcave in i's own strategy. So H lies in Γ^*. A similar but easier argument holds for the auxiliary game in Lemma 3.3(a).

Debreu (1952) introduces abstract economies to facilitate existence proofs of competitive equilibria. They differ from games because the feasible strategies of a player may depend on the strategies chosen by other players: it is a game with feasibility correspondences, one for each player. Peter Sudhölter (private communication) noted that with minor adjustments in the axioms and proofs to address the feasibility constraints, the axiomatization of the Nash equilibrium concept carries over if we replace games in class Γ^* with abstract economies with the same assumptions on strategy spaces and utility functions, and with continuous, convex-valued feasibility correspondences. An axiomatization using *consistency* can be found in Peleg and Sudhölter (1997).

Comparing *independence of identical consequences* and *consistency*. The key axiom that allows comparison of solutions in different games is *independence of identical consequences*. In the earlier axiomatic literature, this role is played by *consistency*. These two axioms are logically independent. We saw in Example 3.4 that the strong Nash equilibrium concept does not satisfy *independence of identical consequences*. It does satisfy *consistency*; see Peleg and Tijs (1996, Thm. 3.2). A strategy $x_i \in X_i$ of player i in a game $G = \langle N, (X_i)_{i\in N}, (u_i)_{i\in N}\rangle$ is *undominated* if there is no other strategy $y_i \in X_i$ with

$$u_i(y_i, x_{-i}) \geq u_i(x_i, x_{-i})$$

for all $x_{-i} \in X_{-i}$, and with strict inequality for some $x_{-i} \in X_{-i}$. Games $G = \langle N, (X_i)_{i\in N}, (u_i)_{i\in N}\rangle$ and $H = \langle N, (X_i)_{i\in N}, (v_i)_{i\in N}\rangle$ satisfying (1) have the same undominated strategies. Therefore, the solution concept that assigns to each finite game G its undominated strategy profiles,

$$\text{sol}(G) = \{x \in X : \text{ for each } i \in N,\ x_i \text{ is undominated}\},$$

satisfies *independence of identical consequences*. But it is not consistent: in the game

	L	R
T	$1,-1$	$-1,1$
B	$-1,1$	$1,-1$

the strategy profile (T, L) is undominated, yet in the one-player reduced game

L	R
-1	1

that arises if the row player commits to T and leaves, L is dominated by R, so L is not a solution of the reduced game, contradicting *consistency*.

On minimax and maximin. We showed that *individual rationality* implies *risk-free back-up*. Therefore, we can replace *risk-free back-up* in Theorem 3.2 by the more demanding *individual rationality* axiom to obtain an alternative axiomatization of the Nash equilibrium concept. But in addition to being less demanding, *risk-free back-up* has another practical advantage: strategies with a constant payoff are easier to recognize than the minimax payoff, which typically requires tedious computations. Analogously, each player i can singlehandedly improve upon utility levels below

$$\sup_{x_i \in X_i} \inf_{x_{-i} \in X_{-i}} u_i(x_i, x_{-i})$$

and we can replace *risk-free back-up* with the stronger requirement that in each solution of the game, players receive at least this maximin utility.

References

R. J. Aumann. Acceptable points in general cooperative n-person games. In A. W. Tucker and R. D. Luce, editors, *Contributions to the Theory of Games IV*, volume 40 of *Annals of Mathematics Studies*, pages 287–324. Princeton University Press, 1959.

L. E. Blume. The statistical mechanics of strategic interaction. *Games and Economic Behavior*, 5: 387–424, 1993.

G. Debreu. A social equilibrium existence theorem. *Proceedings of the National Academy of Sciences*, 38:886–893, 1952.

G. Facchini, F. van Megen, P. Borm, and S. Tijs. Congestion models and weighted potential games. *Theory and Decision*, 42:193–206, 1997.

J. C. Harsanyi and R. Selten. *A General Theory of Equilibrium Selection in Games.* MIT Press, 1988.

J. Milnor. Games against nature. In R. H. Thrall, C. H. Coombs, and R. L. Davis, editors, *Decision Processes*, chapter IV, pages 49–59. John Wiley, 1954.

D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, 14:124–143, 1996.

S. Morris and T. Ui. Best response equivalence. *Games and Economic Behavior*, 49:260–287, 2004.

J. F. Nash. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36:48–49, 1950.

H. Norde, J. Potters, H. Reijnierse, and D. Vermeulen. Equilibrium selection and consistency. *Games and Economic Behavior*, 12:219–225, 1996.

B. Peleg and P. Sudhölter. An axiomatization of Nash equilibria in economic situations. *Games and Economic Behavior*, 18:277–285, 1997.

B. Peleg and S. Tijs. The consistency principle for games in strategic form. *International Journal of Game Theory*, 25:13–34, 1996.

B. Peleg, J. Potters, and S. Tijs. Minimality of consistent solutions for strategic games, in particular for potential games. *Economic Theory*, 7:81–93, 1996.

W. Thomson. On the axiomatic method and its recent applications to game theory and resource allocation. *Social Choice and Welfare*, 18:327–386, 2001.

M. Voorneveld. The possibility of impossible stairways: Tail events and countable player sets. *Games and Economic Behavior*, 68:403–410, 2010.

M. Voorneveld, P. Borm, F. van Megen, S. Tijs, and G. Facchini. Congestion games and potentials reconsidered. *International Game Theory Review*, 1:283–299, 1999.

J. W. Weibull. *Evolutionary Game Theory*. MIT Press, 1995.

W.-T. Wu and J.-H. Jiang. Essential equilibrium points of n-person non-cooperative games. *Scientia Sinica*, 11:1307–1322, 1962.

L. Zhou. The set of Nash equilibria of a supermodular game is a complete lattice. *Games and Economic Behavior*, 7:295–300, 1994.