

Every Normal-Form Game Has a Pareto-Optimal Nonmyopic Equilibrium

Steven J. Brams
Department of Politics
New York University
New York, NY 10012
USA

steven.brams@nyu.edu

Mehmet S. Ismail
Department of Political Economy
King's College London
London WC2R 2LS
UK

mehmet.s.ismail@gmail.com

Abstract

It is well-known that Nash equilibria may not be Pareto-optimal; worse, a *unique* Nash equilibrium may be Pareto-dominated, as in Prisoners' Dilemma. By contrast, we prove a previously conjectured result: Every finite normal-form game of complete information and common knowledge has at least one Pareto-optimal *nonmyopic equilibrium* (NME) in pure strategies, which we define and illustrate. The outcome it gives, which depends on where play starts, may or may not coincide with that given by a Nash equilibrium. We use some simple examples to illustrate properties of NMEs—for instance, that NME outcomes are usually, though not always, maximin—and seem likely to foster cooperation in many games. Other approaches for analyzing farsighted strategic behavior in games are compared with the NME analysis.

1. Introduction

The standard solution concept in noncooperative game theory is that of Nash equilibrium (NE). However, what might be considered a “cooperative outcome” in a significant number of games is not an NE.¹ The best-known example of such a game is Prisoners’ Dilemma.

To justify cooperation in such games, one approach is to posit repeated play of a game. According to the folk theorem of noncooperative game theory, all Pareto-optimal outcomes become NEs if (i) the repetition is infinite or has no definite end and (ii) the players are sufficiently patient. But most real-life games are not played, *de novo*, again and again; moreover, the resulting plethora of NEs in repeated play has little predictive power. Later we discuss other approaches to the dynamic analysis of games.

In this paper, we assume that play starts at an outcome, from which players make farsighted calculations of where play will terminate after a finite series of moves and countermoves *within* a game. This assumption differs radically from the usual assumption of *normal-form games*—represented by a payoff matrix—in which the players are assumed to make independent strategy choices.

By contrast, we specify alternative rules of play that yield *nonmyopic equilibria* (NMEs), which always exist in pure strategies in normal-form games. In 45 of the 57 distinct 2×2 strict ordinal games of conflict with a cooperative outcome, this outcome is an NME, but not necessarily an NE, in all except one.

¹ It is the strategies associated with an outcome, not the outcome itself, that define an NE. Because a unique pair of strategies is associated with each outcome in the 2-person strict ordinal games that we analyze in section 2, for brevity we identify NEs in these games by the outcomes rather than the strategies that yield them.

We begin by defining and illustrating NMEs using Prisoners' Dilemma and other 2-person and n -person games. Although the noncooperative outcome in Prisoners' Dilemma is an NME that coincides with the NE, the news is not all bad: The cooperative outcome is also an NME and *maximin*: It maximizes the minimum payoff that the players receive. (Note that a maximin outcome differs from a maximin strategy outcome, which is the highest payoff a player can guarantee in a game.)

NMEs are based on rules wherein play commences at an outcome, called an initial state, and players can move or countermove from that state according to rules that we specify. If players would not move from an initial state, anticipating all possible moves and countermoves in a game of complete information and common knowledge, then that state is an NME. A state may also be an NME if players would move to it from another state—not just stay at it if they start there—as we discuss later.

In Prisoners' Dilemma, the cooperative outcome is an NME when play commences at it or at one of the two win-lost outcomes. But there are games in which a cooperative outcome, like that in Prisoners' Dilemma, is not an NME. Generally speaking, NMEs are at least middling outcomes for the players, though there are exceptions, which we will illustrate. Unlike NEs, at least one is always Pareto-optimal (our main result, which has been conjectured but not previously proved).

In section 2, we spell out the rules of play and rationality rules for calculating NMEs in 2×2 games, which we apply to Prisoners' Dilemma. In section 3, we use another 2×2 game and two larger games (more strategies, more players) to illustrate other properties of NMEs, such as that they need not be maximin. We then extend the analysis to all finite games and prove that they always contain at least one Pareto-optimal

NME. Our proof, which is very simple, does not identify one but does show that one must exist.

In section 4, we compare our results with those of other theorists who have incorporated dynamic reasoning into the play of games. In section 5, we discuss the normative implications of choosing Pareto-optimal NMEs, especially of cooperation, when they are not myopically stable. This more expansive view for ameliorating conflict offers decision makers an incentive to consider the long-term consequences of their choices, especially when the NEs in a game are not Pareto-optimal or do not exist in pure strategies.

2. Nonmyopic Equilibria: Rules of Play and an Example²

We begin by specifying four rules of play in 2×2 normal-form games of complete information and common knowledge. To describe how NMEs are calculated from each state, we begin by specifying four rules of play:

1. Play starts at an outcome, called the *initial state*, which is at the intersection of the row and column of a payoff matrix.
2. Either player can unilaterally switch its strategy, and thereby change the initial state into a new state, in the same row or column as the initial state. Call the player that switches, who may be either row (R) or column (C), player 1.
3. Player 2 can respond by unilaterally switching its strategy, thereby moving the game to a new state.
4. The alternating responses continue until the player (player 1 or player 2) whose

² This section follows Brams (1994), who gives a full account of the *theory of moves (TOM)*, whose main equilibrium concept is an NME. This account also includes an analysis of three different kinds of power (moving, order, and threat) and several applications; for other applications, see Brams (2011, 2018). Our purpose in this section, after defining NMEs for 2×2 games—the focus of previous studies—is to give general results for larger 2-person and n -person games.

turn it is to move next chooses not to switch its strategy. When this happens, the game terminates in a *final state*, which is the *outcome* of the game.

Note that the sequence of moves and countermoves is strictly alternating: First, say, R moves, then C moves, and so on, until one player stops, at which point the state reached is final and, therefore, the outcome of the game.³

The use of the word “state” is meant to convey the temporary nature of an outcome before players decide to stop switching strategies. We assume that no payoffs accrue to players from being in a state unless it is the final state and, therefore, becomes the outcome (which could be the initial state if the players choose not to move from it).

Rule 1 differs drastically from the corresponding rule of play in standard game theory, in which players simultaneously choose strategies in a normal-form game that determines its outcome. Instead of starting with strategy choices, we assume that players are already in some state at the start of play (the status quo) and receive payoffs from this state only if they choose to stay. Based on these payoffs, they decide, individually, whether or not to change this state in order to do better, which may be physical moves or a thought experiment that anticipates future choices.

In summary, play of a game starts in a state, at which players accrue payoffs only if they remain in that state so that it becomes the outcome of the game. If they do not remain, they still know what payoffs they would have accrued had they stayed; hence, each can make a rational calculation of the advantages of staying versus moving. They move precisely because they calculate that they can do better by switching states, anticipating a better outcome when the move-countermove process comes to rest.

Rules 1-4 say nothing about what causes a game to end but only when: Termination occurs when a “player whose turn it is to move next chooses not to switch its strategy”

³ Rules that allow for backtracking are analyzed in Willson (1998) and applied in Zeager, Ericsson, and Williams (2013).

(rule 4). But when is it rational not to continue moving, or not to move in the first place from the initial state?

To answer this question, we posit a rule of *rational termination*. It prohibits a player from moving from an initial state unless doing so leads to a better (not just the same) final state, based on the following rule:

5. A player will not move from an initial state if this move
 - (i) leads to a less preferred final state (i.e., outcome); or
 - (ii) returns play to the initial state (i.e., makes the initial state the outcome).

We discuss and illustrate shortly how rational players, starting from some initial state, determine by backward induction what the outcome will be.

Condition (i) of rule 5, which precludes moves that result in an inferior state, needs no defense. But condition (ii), which precludes moves that will cause players to cycle back to the initial state, is worth some elaboration. It says that if it is rational for play of a game to cycle back to the initial state after player 1 moves, player 1 will not move in the first place. After all, what is the point of initiating the move-countermove process if play simply returns to “square one,” given that the players receive no payoffs along the way (i.e., before an outcome is reached)?

Not only is there no gain from cycling but, in fact, there may be a loss because of so-called transaction costs—including the psychic energy spent—that players suffer by virtue of making moves that, ultimately, do not change the outcome. Therefore, it seems sensible to assume that player 1 will not trigger a move-countermove process if it only returns the players to the initial state, making it the outcome.

We call rule 5 a *rationality rule*, because it provides the basis for players to determine whether they do better by moving from a state or remaining in it. But another rationality rule is needed to ensure that both players take into account each other’s calculations before deciding to move from the initial state. We call this rule the *two-*

sidedness rule, which assumes that the players have complete information about each other's preferences that is common knowledge, and they act according to the preceding five rules:

6. Each player takes into account the consequences of the other player's rational choices, as well as its own, in deciding whether to move from the initial state or subsequently, based on backward induction. If it is rational for one player to move and the other player not to move from the initial state, then the player who moves takes *precedence*: Its move overrides the player who stays, so the outcome is that induced by the player who moves.⁴

Because players have complete information, they can look ahead and anticipate the consequences of their moves; common knowledge ensures that both players know this. To illustrate the application of the six rules to a 2×2 Prisoners' Dilemma, assume the players' strict ordinal ranking of the four outcomes is as follows: 4 = best; 3 = next best; 2 = next-worst; 1 = worst (see Figure 1).⁵ We next demonstrate, using backward induction, that if (3,3) is the initial state in Prisoners' Dilemma, the players would not move from this state, making it the NME from this state.

Assume R moves first from (3,3), moving play to (4,1), whence play continues cycling counterclockwise back to (3,3), progressing from (4,1) to (2,2) to (1,4) to (3,3). The player (R or C) who makes the next move, shown below each state, alternates:⁶

⁴ If each player prefers to move from an initial state to a different outcome, we consider each outcome to be an NME from that state. But if one player has "order power," it can determine the order of moves and, therefore, that its preferred outcome will be the NME from that state (Brams, 1994).

⁵ Prisoners' Dilemma is game 32 in the classification scheme of Brams (1994, 2011, 2018) for the 57 2×2 conflict games; the other 2×2 game we discuss in this section is game 22, which does not have a name.

⁶ Effectively, this is a game tree, or game in extensive form, showing a sequence of alternating choices of the players, except that instead of branching from top to bottom, as in the usual representation, the choices of the players go sideways, from left to right.

(3,3) [3,3]	(1,4) [3,3]
(4,1) [3,3]	(2,2) [2,2]

Figure 1. Prisoners' Dilemma (Game 32)

Key:

4 = best, 3 = next best, 2 = next worst, 1 = worst

Nash equilibrium (NE) underscored

Nonmyopic equilibria (NMEs) in boldface

NMEs from each initial state in brackets

	State 1 R	State 2 C	State 3 R	State 4 C	State 1
R starts:	(<u>3,3</u>)	→ (4,1)	→ (2,2)	→ (1,4)	→ (<u>3,3</u>)
Survivor:	(3,3)	(2,2)	(2,2)	(1,4)	

The *survivor* is determined by working backwards, after a putative cycle has been completed, which is calculated in the following manner: Assume that the players' alternating moves have taken them counterclockwise from (3,3) eventually to (1,4), at which point C must decide whether to stop at (1,4) or complete the cycle and return to (3,3).

Clearly, C prefers (1,4) to (3,3), so (1,4) is listed as the survivor below (1,4): Because C would *not* move the process back to (3,3) should it reach (1,4), the players know that if the move-countermove process reaches this state, the outcome will be (1,4). We indicate that it is not rational for C to move on from (1,4) by the vertical line

blocking the arrow emanating from (1,4), which we refer to as *blockage*: A player will always stop at a blocked state, wherever it is in the progression.

Would R at the prior state, (2,2), move to (1,4)? Because R prefers (2,2) to the survivor at (1,4)—namely, (1,4)—the answer is no. Once again there is blockage, and (2,2) becomes the new survivor when R must choose between stopping at (2,2) and moving to (1,4).

At the state prior to (2,2), (4,1), C would prefer staying at (2,2) rather than moving to (4,1), so (2,2) again is the survivor if the process reaches (4,1). However, at the initial state, (3,3), because R prefers (3,3) to (2,2), (3,3) becomes the survivor at the initial state, and there is again blockage. But in this case we call the blockage *stoppage*, because it occurs for the *first* time from the initial state; we underscore (3,3) to indicate that it is the last surviving state.

The fact that (3,3) is the survivor at the initial state (3,3) means that it is rational for R not to move from (3,3). That is, after working *backward* from C's choice of completing or not completing the cycle at (1,4), the players can reverse the process and, looking *forward*, determine that it is rational for R not to move from (3,3). Likewise, it is also rational for C not to move from (3,3) because of the symmetry of Prisoners' Dilemma.

An analogous argument shows that it is not rational for either player to move from (2,2), making this outcome an NME from itself, as shown in Figure 1. Applying backward induction from (4,1) and (1,4) indicates that C and R, respectively—each receiving only a payoff of 1—would move to (2,2), where play would stop. But, in fact, we amend this calculation of NMEs from (4,1) and (1,4) according to the following convention:

Two-sidedness convention (TSC): If one player (say, C), by moving, can induce a better state for itself than by staying—but R by moving can induce a state

Pareto-superior to C's induced state—then R will move, even if it otherwise would prefer to stay, to effect a better outcome.

To illustrate this convention, observe that from (4,1), C can induce a better state for itself by moving to (2,2). But because R, by moving first to (3,3), can induce a state Pareto-superior to (2,2), it behooves R *not* to stay at (4,1) but instead to move to (3,3). Moreover, it is also in C's interest to defer its move to (2,2) to enable R to implement (3,3).

Although we could make TSC a new rule (i.e., rule 7), it seems better to call it a “convention,” because it clarifies a circumstance when rule 6 is operative—that is, when a move by a player takes precedence. In our example, although C at (4,1) immediately benefits by moving to (2,2) rather than staying at (4,1), TSC says that it nevertheless is rational for C to stay in order to allow R to move first to (3,3).

3. Cooperative Outcomes and the Pareto-Optimality of at Least One NME

Define a *cooperative outcome* in a 2×2 strict ordinal game to be (3,3), as in Prisoners' Dilemma, or (3,4) or (4,3), all of which are a subset of maximin outcomes. Of the 57 distinct 2×2 *conflict games* that contain no mutually best (4,4) outcome, 45 (79 percent) have either one or two cooperative outcomes; the maximin outcomes in the other 12 conflict games are (2,4), (4,2), (2,3), or (3,2).⁷

The cooperative outcomes in the 45 games in which each player receives at least its next-best payoff of 3 are, with one exception (game 22 in the appendices of Brams, 1994, 2011), NMEs. This exception is shown in Figure 2.

⁷ In addition, there are 21 distinct no-conflict games with a mutually best (4,4) outcome—all of which are NMEs—giving a total of 78 distinct 2×2 strict ordinal games.

(3,3)	(2,4)
(4,1)	(1,2)

Figure 2. Game 22

Key:

4 = best, 3 = next best, 2 = next worst, 1 = worst

Nash equilibrium (NE) underscored

Nonmyopic equilibrium (NME) in boldface, from every initial state

In this game, (2,4) is the unique NME, wherever play starts. Although there is a cooperative outcome, (3,3), it is not an NME: Backward induction from (3,3) shows that C would have an immediate incentive to move to (2,4), where play would stop; R would follow a more circuitous route to (2,4), passing first through (4,1) and then (1,2), before arriving at (2,4). Unlike Prisoners' Dilemma, in which both (3,3) and (2,2) are NMEs, (2,4) is the only NME in game 22.

In games larger than 2×2 , define a *cooperative outcome* to be one that maximizes the minimum payoff to the players, which is true of (3,3) in game 22.⁸ Like game 22, cooperative outcomes in larger games always exist but may not be NMEs. As a case in point, consider the 2×3 game in Figure 3, in which the two players rank the six outcomes from best (6) to worst (1).

⁸ Thus, in larger game, cooperative outcomes coincide with maximin outcomes and are not, as in 2×2 games, a subset of them. This definition simplifies the analysis of these games, but it might be refined in future work.

(6,1)	(4,4)	(1,6)
(5,2)	(3,3)	(2,5)

Figure 3. A 2 × 3 Game

Key:

6 = best, 5 = next best, . . . , 1 = worst

Nash equilibrium (NEs) underscored

Nonmyopic equilibrium (NME) in boldface, from every initial state

In this game, (4,4) is not an NME, because C would move to (1,6), whence R would move to (2,5), from which C would not move, because subsequent rational moves would return play to (4,4), where C would do worse. By similar reasoning, it is not difficult to show that starting from any outcome other than (4,4), the players would stay at or move to (2,5), making it the unique NME in the 2 × 3 game (it is also the unique NE).⁹

This game illustrates that a maximin outcome like (4,4)—or (3,3) in game 22—which one might expect to be nonmyopically stable, will not always be so. Instead, a more one-sided outcome like (2,5)—or (2,4) in game 22—may be the only NME. Like (2,4) in game 22, (2,5) in the 2 × 3 game is associated with the dominant strategy of one player (C) and is the unique Nash equilibrium in this game.

Other outcomes one might not expect to be NMEs include (2,2) in Prisoners' Dilemma. This outcome is Pareto-dominated by the cooperative outcome, (3,3), which is

⁹ In Brams (1994, pp. 11-17), rules of play for larger games are illustrated by a 3 × 3 game.

also an NME—not only from itself but also from (4,1) and (1,4)—whereas (2,2) is an NME only if play starts at this state.

In fact, (2,2) in Prisoners' Dilemma is the only Pareto-dominated NME in all 57 of the 2×2 conflict games, suggesting that Pareto-dominated NMEs in larger games are likely to be rare. Fortunately, in both 2×2 and larger games, it is never the case that all NMEs are Pareto-dominated if the preferences of players are strict, as we prove next.

Theorem. *Assume that the rules of play in any 2-person or n-person normal-form game preclude players from repeating a set of moves that returns play to a previously visited state.¹⁰ In every finite normal-form game of complete information and common knowledge, at least one NME must be Pareto-optimal.*

Proof. Assume the contrary—that all NMEs are Pareto-dominated. Consider an NME from some outcome, and assume it is Pareto-dominated by some Pareto-optimal outcome X (there must be at least one X).

Consider all backward-induction paths from X back to itself. Because of the prohibition on returning to a previously visited state, both the number and the length of paths in the backward-induction process must be finite. For every Pareto-dominated NME, there will be an X that Pareto-dominates it and so will displace it along a path from X back to itself, thereby making X an NME from itself or go into another Pareto-optimal NME.

¹⁰ This assumption automatically holds in 2×2 games, because there is only one path for each player, cycling either clockwise or counterclockwise back to the initial state.

Therefore, not all NMEs can be Pareto-dominated, as assumed. Instead, some X must be an NME from itself or go into one. Because there is at least one X , there must be at least one Pareto-optimal NME. Q.E.D.

To illustrate Theorem, recall that $(2,2)$ is a Pareto-dominated NME from itself in Prisoners' Dilemma. But starting at $(3,3)$, which Pareto-dominates $(2,2)$, $(3,3)$ displaces $(2,2)$ along a path back to itself, rendering $(3,3)$ the NME from itself. Starting at $(4,1)$ or $(1,4)$, these outcomes go into $(3,3)$ by TSC, as we showed earlier.

As an illustration of Theorem in an n -person game, consider the 3-person 2×2 game in Figure 4. In this game, in which the third player chooses the left or right matrix and receives the third payoff at every state, $(5,6,5)$, shown in boldface, is the unique NME in this game. Note, however, that $(5,6,5)$ is not an NE—the Pareto-dominated $(4,5,4)$, which is underscored, is the unique NE in the 3-person game.

To demonstrate that $(5,6,5)$ is the unique NME in this game, it is not difficult to show that the NMEs (and also NEs) in the left and right 2×2 matrices are $(3,7,3)$ and $(4,5,4)$, respectively, if we treat each matrix as a distinct 2×2 game and ignore the payoffs of the third player. Because the third player prefers $(4,5,4)$ to $(3,7,3)$, this is the unique NE in the 3-person game—each player would do worse if it departed from it.

However, $(5,6,5)$ in the left matrix Pareto-dominates $(4,5,4)$ in the right matrix, so it is not rational for the matrix player to move from $(5,6,5)$ to $(7,3,6)$ in the right matrix—benefitting immediately—because play will then terminate at $(4,5,4)$ in the right matrix. Although both R and C can do immediately better by moving from $(5,6,5)$ if play stays in the left matrix, the NME for these players in this matrix is $(3,7,3)$.

$(3,7,3)$	$(5,6,5)$	<u>$(4,5,4)$</u>	$(7,3,6)$
-----------	-----------------------------	-----------------------------	-----------

(2,4,1)	(8,1,2)	(1,8,7)	(6,2,8)

Figure 4. A 2 × 2 × 2 Game

Key:

8 = best, 7 = next best, . . . , 1 = worst

Nash equilibrium (NE) underscored

Nonmyopic equilibrium (NME) in boldface, from every initial state

Once at (3,7,3), however, the matrix player will move to (4,5,4), where play will terminate. Like Prisoners' Dilemma, this is an example of a game in which the unique NME, (5,6,5), Pareto-dominates the unique NE, (4,5,4), illustrating the advantage to all three players of thinking ahead rather than myopically.¹¹

Theorem does not show how to find NMEs in normal-form games larger than 2 × 2, only that at least one must be Pareto-optimal. While the task of finding NMEs in 2 × 2 games is straightforward, the determination of NMEs in larger games, such as those in Figures 3 and 4, is less straightforward, because there are in general multiple paths back to an initial state that need to be checked, and the results of backward induction from each compared, to determine which outcomes are NMEs that farsighted players will choose.

¹¹ All our examples assume that the preferences of the players are strict. If they are not, more than one NME may survive the backward induction process because of ties. But ties have no effect on Theorem—at least one NME must be Pareto-optimal in the sense of there being no other outcome that is weakly preferred to it by the players.

We bypassed a systematic NME analysis of the games in Figures 3 and 4 to argue less formally that one outcome in each game is the unique NME, wherever play starts. We have not tried to provide an algorithm that will find all NMEs in larger 2-person and all n -person games. What Theorem does provide is assurance that at least one Pareto-optimal NME exists, starting from some initial state.

In the case of the 2×2 games, it is worth reviewing our findings:

1. If there is a cooperative outcome that is at least next best for both players (true in 45 of the 57 conflict games, including Prisoners' Dilemma), almost invariably (exception: game 22) this outcome will be an NME and maximin.
2. Prisoners' Dilemma has a Pareto-dominated NME, (2,2), if play starts there, that coincides with its unique NE.

Although several other 2×2 games have unique Pareto-optimal NMEs that Pareto-dominate a unique Pareto-dominated NE, Prisoners' Dilemma is the only one in which the Pareto-dominated NE is also an NME, in part explaining its prominence in the game-theoretic literature.

4. Other Approaches to the Dynamic Analysis of Games

We, of course, are not the first to assume that players look ahead in formulating their rational choices. In the Introduction, we alluded to the repeated play of games as a model for showing how cooperation can emerge in games, like Prisoners' Dilemma, when repeated. We suggested, however, that when a game is simply repeated, it may not realistically describe the choices of players if it is rational for them to move, or respond to moves, within the game—not simply play a new game.

Evolutionary game theory offers an alternative view of the dynamics of strategy choices, because it takes into account changes in the population of players over time as some types grow and others decline. It has been most often applied in biology to the competition within a species whose members have different traits and choose different strategies. It is less relevant to explaining the choices of human players and, in particular, their rational calculations that underlie evolutionarily stable strategies, the main equilibrium concept in this literature.

Preceding Brams's (1994, 2011, 2018) work on theory of moves, Brams and Wittman (1981) and Kilgour (1984) proposed move-based calculations in normal-form games of complete information and common knowledge, but they did not provide general results for such games or subclasses of them. Even earlier, theorists had studied how the iterated elimination of strictly dominated strategies can reduce games to smaller games and even a single outcome.

This approach assumes that players can apprehend the effects of successive strategy eliminations, both of themselves and by other players, before applying the same reasoning to the reduced versions of the original game. This can be demanding but has some applications—for example, in the determination of “sophisticated” strategies in voting games (Farquharson, 1969). A related notion of dynamic reasoning within a game is rationalizability, as formulated by Bernheim (1984) and Pearce (1984).

In cooperative game theory, n -person coalitional games have been studied since von Neumann and Morgenstern (1944); the core, from which no player or coalition of players finds it profitable to deviate, is a central concept in this literature. Von Neumann

and Morgenstern proposed a more sophisticated concept of stability, the stable set, that is internally and externally stable, but it does not always exist.

The core and stable set are myopic solution concepts, because possible counterdeviations by players to a deviation are not considered. This is not true of the bargaining set (Aumann and Maschler, 1964), which allows for objections and counterobjections and, therefore, offers a more dynamic view of how stability can be achieved.

More recently, so-called expectation functions have been incorporated into coalitional games that specify possible moves among states that lead to farsighted stability (see, e.g., Dutta and Vohra, 2017, and references therein). Applied to a 2-person normal-form game, they might specify, as we have, that the order of moves strictly alternates and that players not repeat a prior set of moves before play returns to the initial state (as pointed out in note 10, such repetition cannot occur in 2×2 games).

States need not be the outcomes of a normal-form game, and they may be changed by coalitions of players.¹² Whether stationary points, like NMEs, are Pareto-optimal or maximin are not standard questions in this literature.

Recursive games of perfect information are a subclass of stochastic games in which there are finitely many states and each player becomes “active” at some state (see, Flesch et al., 2010, and references therein).¹³ The active player can choose to (i) quit, in which case the play goes into an absorbing state, which is similar to a stationary point,

¹² More recently, Karos and Kasper (2018) extend Dutta and Vohra’s (2017) notion of an expectation function to allow for coalitions to make moves based on counterfactual reasoning, such as that incorporated in the two-sidedness convention (TSC) whereby players counter the expectation that their moves will be based solely on backward induction.

¹³ Stochastic games were first introduced by Shapley (1953). Kuipers et al. (2016) provide an overview of the more recent literature.

and positive payoffs are collected; or (ii) move to another state, after which some other player becomes active.

Unlike the normal-form games we assume, in recursive games (i) there can be infinitely many states; (ii) the utility functions of players may be discounted or may be defined by the expected average reward; and (iii) the transition from one state to another may be either stochastic or deterministic. A related class of games is Dynkin games, in which some player is active at every state and has only two actions: to quit or to continue play, which makes another player active (see Solan and Vieille, 2003, and references therein).

We have touched upon different approaches that have been used to model the dynamics of player and coalitional choices in games; there are certainly more. In the concluding section, we review our results, discussing some of the advantages of NMEs, especially compared with NEs, in explaining strategic behavior that may, speaking normatively, help players achieve cooperative outcomes in games.

5. Conclusions

We have assumed a specific game form (normal form) and, at the outset, described rules of play in 2×2 strict ordinal conflict games that define NMEs from every initial state. Backward induction yields one or more NMEs, and at least one is Pareto-optimal in both 2-person and n -person finite games of complete information and common knowledge, which we illustrated with several examples.

We did not provide a general algorithm for calculating NMEs in games larger than 2×2 , but our proof of the Pareto-optimality of at least one NME does not depend on such an algorithm as long as play always returns to the initial state—by whatever path

the players choose—and no previous set of moves is repeated, which ensures that the number of moves is finite. In 45 of the 57 2×2 conflict games (79 percent), at least one NME is a cooperative outcome, at which each player receives at least its next-best payoff. In only one of these games (Prisoners' Dilemma) is there an NME, (2,2), that is not Pareto-optimal but, as required by Theorem, Prisoners' Dilemma has a second NME, (3,3), that is Pareto-optimal.

This is not to claim that rational players will always attain such an NME. In Prisoners' Dilemma, for example, if they start at (2,2), they will be stuck there, though if they start at any of the other three outcomes, (3,3) will be the NME. Perhaps surprisingly, there is one 2×2 game (game 22) that, wherever play starts, the unique NME is (2,4) and not the cooperative outcome, (3,3), in this game.

Prisoner's Dilemma and game 22 are decidedly the exceptions among the 45 2×2 games with cooperative outcomes. For the 12 other 2×2 conflict games, the NMEs are maximin outcomes. It follows that game 22 is the only 2×2 conflict game in which there is no maximin NME—(3,3) is the maximin outcome in this game—but, as required by Theorem, the (2,4) NME is Pareto-optimal.

The picture is very different for NEs.¹⁴ There are two games (26 and 27) in which the unique pure-strategy NE is Pareto-dominated, whereas the outcome that Pareto-dominates it is the unique cooperative NME. In another three games (29-31), there is no pure-strategy NE, but a unique cooperative outcome is the NME. In still six more games (42-47), there is also no pure-strategy NE, but each of these games has two Pareto-

¹⁴ We do not show the games we list in this paragraph, but they can be found in the appendices of Brams (1994, 2011, 2018).

optimal NMEs—one of which is a cooperative outcome in two of these games (46 and 47)—whose choice depends on where play starts.

Altogether, there are nine games in which the NE is Pareto-dominated or in mixed strategies, but each of these games has one or more Pareto-optimal NMEs, in five of which the outcomes are cooperative. In effect, NMEs offer an escape from both Pareto-dominated NEs or the ambiguity of mixed strategies, at least when players can associate cardinal utilities with the outcomes and can therefore make expected-value calculations.

Multiple NEs in many games, even when the outcomes they yield are cooperative, also pose difficult choices for players. NMEs can be helpful in singling out one of these NEs, depending on where play commences.

The farsighted thinking that undergirds NMEs does not solve all problems, even in 2×2 games, of ensuring a cooperative outcome if one exists, as we saw in both Prisoners' Dilemma and game 22. But the advantages to the players of looking ahead, knowing that at least one NME is Pareto-optimal and frequently maximin, seem considerable in 2×2 as well as larger games. However, the NME calculations in larger games, as we illustrated in the case of the Figure 3 and 4 games, are more arduous and so less intuitive for decision makers to realize.

All our results depend on modeling moves as sequential, without specifying any particular order unless the players have a common interest in who, if anybody, moves first and follows thereafter. Because moves may take the form of mental calculations—not necessarily physical actions—players may appear to make simultaneous choices when in fact they earlier played out in their minds different scenarios about what move to

make and what their reaction would be to the moves of other players. The give-and-take of negotiations among two or more players may also be modeled by the NME analysis.

To be sure, we may not always observe the moves and countermoves of players in strategic situations requiring them to think ahead or bargain with other players. But we think that nonmyopic strategic behavior can and, speaking normatively, should be modeled, especially in light of our finding that it leads to Pareto-optimal outcomes that are frequently cooperative in nature.

References

- Aumann, R. J., and M. Maschler (1964). "The Bargaining Set for Cooperative Games."
 In M. Dresher, L. S. Shapley, and A. W. Tucker (eds.) (1964). *Advances in Game Theory*, pp. 443-476. Princeton, NJ: Princeton University Press.
- Bernheim, B. Douglas (1984). "Rationalizable Strategic Behavior." *Econometrica* 52, no. 4 (July): 1007-1028.
- Brams, Steven J. (1994). *Theory of Moves*. Cambridge, UK: Cambridge University Press.
- Brams, Steven J. (2011). *Game Theory and the Humanities: Bridging Two Worlds*. Cambridge, MA: MIT Press.
- Brams, Steven J. (2018). *Divine Games: Game Theory and the Undecidability of a Superior Being*. Cambridge, MA: MIT Press.
- Brams, Steven J., and Donald Wittman (1981). "Nonmyopic Equilibria in 2 × 2 Games." *Conflict Management and Peace Science* 6, no. 1 (Fall): 39-62.
- Dutta, Bhaskar, and Rajiv Vohra (2017). "Rational Expectations and Farsighted Stability." *Theoretical Economics* 12, no. 3 (September): 119-1227.
- Farquharson, Robin (1969). *Theory of Voting*. New Haven, CT: Yale University Press.
- Flesch, J., Kuipers, J., Schoenmakers, G., and Vrieze, K. (2010). "Subgame Perfection in Positive Recursive Games with Perfect Information." *Mathematics of Operations Research* 35, no. 1 (February): 193-207.
- Karos, D., and L. Kasper (2018). "Farsighted Rationality." Preprint, Maastricht University. <https://cris.maastrichtuniversity.nl/ws/files/25765130/RM18011.pdf>
- Kilgour, D. Marc (1984). "Equilibria for Farsighted Players." *Theory and Decision* 16,

- no. 2 (March): 135-157.
- Kuipers, J., Flesch, J., Schoenmakers, G., and K. Vrieze (2016). "Subgame-Perfection in Recursive Perfect Information Games, where Each Player Controls One State." *International Journal of Game Theory* 45, nos. 1-2 (October): 205-237.
- Pearce, David B. (1984). "Rationalizable Strategic Behavior and the Problem of Perfection." *Econometrica* 52, no. 4 (July): 1029-1050.
- Shapley, L. S. (1953). "Stochastic Games." *Proceedings of the National Academy of Sciences* 39, no. 10: 1095-1100.
- Solan, Eilon, and Nicholas Vieille (2003). "Deterministic Multi-Player Dynkin Games." *Journal of Mathematical Economics* 39, no. 8 (November): 911-929.
- Von Neumann, John, and Oskar Morgenstern (1944). *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press.
- Willson, Stephen J. (1998). "Long-Term Behavior in the Theory of Moves." *Theory and Decision* 45, no. 3: 167-180.
- Zeager, Lester R., Richard E. Ericson, and John H. P. Williams (2013). *Refugee Negotiations from a Game-Theoretic Perspective*. Dordrecht, The Netherlands: Republic of Letters.