



Why Winning in Rock-Paper-Scissors (and in Life) Isn't Everything

What does John Nash's game theory equilibrium concept look like in Rock-Paper-Scissors?

By Patrick Honner

Rock-Paper-Scissors works great for deciding who has to take out the garbage. But have you ever noticed what happens when, instead of playing best of three, you just let the game continue round after round? At first, you play a pattern that gives you the upper hand, but then your opponent quickly catches on and turns things in her favor. As strategies evolve, a point is reached where neither side seems to be able to improve any further. Why does that happen?

In 1950, the mathematician John Nash proved that in any kind of game with a finite number of players and a finite number of options — like Rock-Paper-Scissors — a mix of strategies always exists where no single player can do any better by changing their own strategy alone. The theory behind such stable strategy profiles, which came to be known as "[Nash equilibria](#)," revolutionized the field of game theory, altering the course of economics and changing the way everything from political treaties to network traffic is studied and analyzed. And it earned Nash the [Nobel Prize in 1994](#).

So, what does a Nash equilibrium look like in Rock-Paper-Scissors? Let's model the situation with you (Player A) and your opponent (Player B) playing the game over and over. Each round, the winner earns a point, the loser loses a point, and ties count as zero.

Now, suppose Player B adopts the (silly) strategy of choosing Paper every turn. After a few rounds of winning, losing, and tying, you are likely to notice the pattern and adopt a winning counterstrategy by choosing Scissors every turn. Let's call this strategy profile (Scissors, Paper). If every round unfolds as Scissors vs. Paper, you'll slice your way to a perfect record.

But Player B soon sees the folly in this strategy profile. Observing your reliance on Scissors, she switches to the strategy of always choosing Rock. This strategy profile (Scissors, Rock) starts winning for Player B. But of course, you now switch to Paper. During these stretches of the game, Players A and B are employing what are known as "pure" strategies — a single strategy that is chosen and repeatedly executed.

Clearly, no equilibrium will be achieved here: For any pure strategy, like "always choose Rock," a counterstrategy can be adopted, like "always choose Paper," which will force another change in strategy. You and your opponent will forever be chasing each other around the circle of strategies.

But you can also try a "mixed" strategy. Let's assume that, instead of choosing only one strategy to play, you can randomly choose one of the pure strategies each round. Instead of "always play Rock,"

a mixed strategy could be to “play Rock half the time and Scissors the other half.” Nash proved that, when such mixed strategies are allowed, every game like this must have at least one equilibrium point. Let’s find it.

So, what’s a sensible mixed strategy for Rock-Paper-Scissors? A reasonable intuition would be “choose Rock, Paper or Scissors with equal probability,” denoted as $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. This means Rock, Paper and Scissors are each chosen with probability $\frac{1}{3}$. Is this a good strategy?

Well, suppose your opponent’s strategy is “always choose Rock,” a pure strategy that can be represented as (1,0,0). How will the game play out under the strategy profile $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ for A and (1,0,0) for B?

In order to get better picture of our game, we’ll construct a table that shows the probability of each of the nine possible outcomes every round: Rock for A, Rock for B; Rock for A, Paper for B; and so on. In the chart below, the top row indicates Player B’s choice, and the leftmost column indicates Player A’s choice.

A B	R	P	S
R	$\frac{1}{3}$	0	0
P	$\frac{1}{3}$	0	0
S	$\frac{1}{3}$	0	0

Each entry in the table shows the probability that the given pair of choices is made in any given round. This is simply the product of the probabilities that each player makes their respective choice. For example, the probability of Player A choosing Paper is $\frac{1}{3}$, and the probability of player B choosing Rock is 1, so the probability of (Paper for A, Rock for B) is $\frac{1}{3} \times 1 = \frac{1}{3}$. But the probability of (Paper for A, Scissors for B) is $\frac{1}{3} \times 0 = 0$, since there is a zero probability of Player B choosing Scissors.

So how does Player A fare in this strategy profile? Player A will win one-third of the time (Paper, Rock), lose one-third of the time (Scissors, Rock) and tie one-third of the time (Rock, Rock). We can compute the number of points that Player A will earn, on average, each round by computing the sum of the product of each outcome with its respective probability:

$$\frac{1}{3}(1) + \frac{1}{3}(0) + \frac{1}{3}(-1) = 0$$

This says that, on average, Player A will earn 0 points per round. You will win, lose and tie with equal likelihood. On average, the number of wins and losses will even out, and both players are essentially headed for a draw.

But as we’ve already discussed, you can do better by changing your strategy, assuming your opponent doesn’t change theirs. If you switch to the strategy (0,1,0) (“choose Paper every time”), the probability chart will look like this

A B	R	P	S
R	0	0	0
P	1	0	0
S	0	0	0

Each time you play, your Paper will wrap your opponent’s Rock, and you’ll earn one point every

round.

So, this pair of strategies — $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ for A and $(1, 0, 0)$ for B — is not a Nash equilibrium: You, as Player A, can improve your results by changing your strategy.

As we’ve seen, pure strategies don’t seem to lead to equilibrium. But what if your opponent tries a mixed strategy, like $(\frac{1}{2}, \frac{1}{4}, \frac{1}{4})$? This is the strategy “Rock half the time; Paper and Scissors each one quarter of the time.” Here’s the associated probability chart:

A B	R	P	S
R	$\frac{1}{6}$	$\frac{1}{12}$	$\frac{1}{12}$
P	$\frac{1}{6}$	$\frac{1}{12}$	$\frac{1}{12}$
S	$\frac{1}{6}$	$\frac{1}{12}$	$\frac{1}{12}$

Now, here’s the “payoff” chart, from Player A’s perspective; this is the number of points Player A receives for each outcome.

A B	R	P	S
R	0	-1	1
P	1	0	-1
S	-1	1	0

We put the two charts together, using multiplication, to compute how many points, on average, Player A will earn each round.

$$\frac{1}{6}(0) + \frac{1}{12}(-1) + \frac{1}{12}(1) + \frac{1}{6}(1) + \frac{1}{12}(0) + \frac{1}{12}(-1) + \frac{1}{6}(-1) + \frac{1}{12}(1) + \frac{1}{12}(0) = 0$$

On average, Player A is again earning 0 points per round. Like before, this strategy profile, $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ for A and $(\frac{1}{2}, \frac{1}{4}, \frac{1}{4})$ for B, ends up in a draw.

But also like before, you as Player A can improve your results by switching strategies: Against Player B’s $(\frac{1}{2}, \frac{1}{4}, \frac{1}{4})$, Player A should play $(\frac{1}{4}, \frac{1}{2}, \frac{1}{4})$. This has a probability chart of

A B	R	P	S
R	$\frac{1}{8}$	$\frac{1}{16}$	$\frac{1}{16}$
P	$\frac{1}{4}$	$\frac{1}{8}$	$\frac{1}{8}$
S	$\frac{1}{8}$	$\frac{1}{16}$	$\frac{1}{16}$

and this net result for A:

$$\frac{1}{8}(0) + \frac{1}{16}(-1) + \frac{1}{16}(1) + \frac{1}{4}(1) + \frac{1}{8}(0) + \frac{1}{8}(-1) + \frac{1}{16}(-1) + \frac{1}{16}(1) + \frac{1}{16}(0) = \frac{1}{16}$$

That is, under this strategy profile — $(\frac{1}{4}, \frac{1}{2}, \frac{1}{4})$ for A and $(\frac{1}{2}, \frac{1}{4}, \frac{1}{4})$ for B — Player A nets $\frac{1}{16}$ of a point per round on average. After 100 games, Player A will be up 6.25 points. There’s a big incentive for Player A to switch strategies. So, the strategy profile of $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ for A and $(\frac{1}{2}, \frac{1}{4}, \frac{1}{4})$ for B is not a

Nash equilibrium, either.

But now let's consider the pair of strategies $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ for A and $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ for B. Here's the corresponding probability chart:

A B	R	P	S
R	$\frac{1}{9}$	$\frac{1}{9}$	$\frac{1}{9}$
P	$\frac{1}{9}$	$\frac{1}{9}$	$\frac{1}{9}$
S	$\frac{1}{9}$	$\frac{1}{9}$	$\frac{1}{9}$

Symmetry makes quick work of the net result calculation:

$$\frac{1}{9}(0) + \frac{1}{9}(-1) + \frac{1}{9}(1) + \frac{1}{9}(1) + \frac{1}{9}(0) + \frac{1}{9}(-1) + \frac{1}{9}(-1) + \frac{1}{9}(1) + \frac{1}{9}(0) = 0$$

Again, you and your opponent are playing to draw. But the difference here is that no player has an incentive to change strategies! If Player B were to switch to any imbalanced strategy where one option — say, Rock — were played more than the others, Player A would simply alter their strategy to play Paper more frequently. This would ultimately yield a positive net result for Player A each round. This is precisely what happened when Player A adopted the strategy $(\frac{1}{4}, \frac{1}{2}, \frac{1}{4})$ against Player B's $(\frac{1}{2}, \frac{1}{4}, \frac{1}{4})$ strategy above.

Of course, if Player A switched from $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ to an imbalanced strategy, Player B could take advantage in a similar manner. So, neither player can improve their results solely by changing their own individual strategy. The game has reached a Nash equilibrium.

The fact that all such games have such equilibria, as Nash proved, is important for several reasons. One of those reasons is that many real-life situations can be modeled as games. Whenever a group of individuals is caught in the tension between personal gain and collective satisfaction — like in a negotiation, or a competition for shared resources — you'll find strategies being employed and payoffs being evaluated. The ubiquitous nature of this mathematical model is part of the reason Nash's work has been so impactful.

Another reason is that a Nash equilibrium is, in some sense, a positive outcome for all players. When reached, no individual can do better by changing their own strategy. There might exist better collective outcomes that could be reached if all players acted in perfect cooperation, but if all you can control is yourself, ending up at a Nash equilibrium is the best you can reasonably hope to do.

And so, we might hope that "games" like economic incentive packages, tax codes, treaty parameters and network designs will end in Nash equilibria, where individuals, acting in their own interest, all end up with something to be happy about, and systems are stable. But when playing these games, is it reasonable to assume that players will naturally arrive at a Nash equilibrium?

It's tempting to think so. In our Rock-Paper-Scissors game, we might have guessed right away that neither player could do better than playing completely randomly. But that's in part because all player preferences are known to all other players: Everyone knows how much everyone else wins and loses for each outcome. But what if preferences were secret and more complex?

Imagine a new game in which Player B scores three points when she defeats Scissors, and one point

for any other victory. This would alter the mixed strategy: Player B would play Rock more often, hoping for the triple payoff when Player A chooses Scissors. And while the difference in points wouldn't directly affect Player A's payoffs, the resulting change in Player B's strategy would trigger a new counterstrategy for A.

And if every one of Player B's payoffs was different, and secret, it would take some time for Player A to figure out what Player B's strategy was. Many rounds would pass before Player A could get a sense of, say, how often Player B was choosing Rock, in order to figure out how often to choose Paper.

Now imagine there are 100 people playing Rock-Paper-Scissors, each with a different set of secret payoffs, each depending on how many of their 99 opponents they defeat using Rock, Paper or Scissors. How long would it take to calculate just the right frequency of Rock, Paper or Scissors you should play in order to reach an equilibrium point? Probably a long time. Maybe longer than the game will go on. Maybe even longer than the lifetime of the universe!

At the very least, it's not obvious that even perfectly rational and reflective players, playing good strategies and acting in their own best interests, will end up at equilibrium in this game. This idea lies at the heart of a [paper posted online in 2016](#) that proves there is no uniform approach that, in all games, would lead players to even an approximate Nash equilibrium. This is not to say that perfect players never tend toward equilibrium in games — they often do. It just means that there's no reason to believe that just because a game is being played by perfect players, equilibrium will be achieved.

When we design a transportation network, we might hope that the players in the game, travelers each seeking the fastest way home, will collectively achieve an equilibrium where nothing is gained by taking a different route. We might hope that the invisible hand of John Nash will guide them so that their competing and cooperating interests — to take the shortest possible route yet avoid creating traffic jams — produce an equilibrium.

But our increasingly complex game of Rock-Paper-Scissors shows why such hopes may be misplaced. The invisible hand may guide some games, but others may resist its hold, trapping players in a never-ending competition for gains forever just out of reach.

Exercises

1. Suppose Player B plays the mixed strategy $(\frac{1}{2}, \frac{1}{2}, 0)$. What mixed strategy should A play to maximize wins in the long run?
2. Suppose Player B plays the mixed strategy $(\frac{1}{6}, \frac{2}{6}, \frac{3}{6})$. What mixed strategy should A play to maximize wins in the long run?
3. How might the dynamics of the game change if both players are awarded a point for a tie?

Correction: This column was revised on April 3, 2018, to reflect that $1/4 \times 1/3 = 1/12$. The results for one set of strategy profiles was also corrected to indicate that Paper would be played against Rock.

This article was reprinted on [Wired.com](#).