

# At Lunch with Freeman Dyson

William Press

**I**N THIS ESSAY, I would like to tell the story of a minor discovery in mathematical game theory that Freeman Dyson and I made in 2011. Dyson was a personal friend and one of the great mathematical physicists of the twentieth century. He died in 2020, at the age of ninety-six. He was famously self-effacing, which is not to say that he lacked an accurate opinion of his own abilities. Freeman would deny that he had done anything at all and then allow friends—or even strangers—to vehemently contradict him. Our discovery was not of that character. It really *was* very minor. The reasons for telling the story now are less about the discovery itself and more about the tendency of scientists to seek lessons in moral philosophy in the least likely of places—high-school algebra, for example.

Imagine that a group of scientists gather to play a kind of terror game. They must propose scenarios that, should they eventuate, would shake their belief in the foundations of their fields. The mathematician’s proposed terror is that a long message, in English, is found to be encoded—in excess of any plausible random probability—somewhere in the first billion digits of pi.<sup>1</sup> The physicist’s terror is that the interaction cross-section of a fundamental particle will have significantly different values when measured in different places on earth, or in the same place at different times.<sup>2</sup> The biologist’s terror is that some feature of the living world will be unexplainable by the principle of natural selection. Within biology’s subspecialty of evolution theory, there is a small area of study known as *evolution of cooperation*. That study, some would say, lies closest to the biologist’s terror. That makes it worth poking at.

**I**N BIOLOGY, a cooperator is an individual who pays a cost for another individual to receive a benefit. When cooperation is mutually beneficial to two individuals of the same or different species—a condition termed *direct reciprocity*—then it is favored by natural selection. There are other possibilities. In so-called kin selection, an individual’s self-sacrifice may be favored if, on average, it helps another individual in the same gene pool to survive.<sup>3</sup> The unit of survival is understood in this case to be not the individual, but the gene that two individuals share.<sup>4</sup> It is harder to understand why individuals cooperate when

defection would be more favorable or when the reciprocity is only indirect.

Suppose that two microbe species, *A* and *B*, both need processed nutrients *a* and *b*. The cooperative state might be that *A* produces *a*, *B* produces *b*, and each secretes a portion of its nutrient for the benefit of the other. But this equilibrium is not evolutionarily stable: a defecting *A* with a mutation that halts its sharing of *a* becomes a free rider, benefitting from *B* without paying the fare. Free riders, avoiding a cost, will tend to take over the population. The evolutionarily stable endpoint is noncooperation, even though cooperation would be better for both species.

Cooperation among humans seems hardest of all to understand. “Humans are the champions of cooperation,” Martin Nowak has remarked. “From hunter-gatherer societies to nation-states, cooperation is the decisive organizing principle of human society.”<sup>5</sup> In much, if not most, of our cooperation, reciprocity is indirect. To be sure, some people give money to universities in the hope of getting their own children admitted—kin selection—but many more give to charities that are of no direct benefit to themselves or their kin. Many billionaires become philanthropists, but from the standpoint of evolution theory, why is this? A quirk of our culture, maybe? But cultures, too, compete for dominance with other contemporaneous cultures, and by a process akin to natural selection. Are we to understand that generosity is selectively favored? Or are the generous billionaires only transient?

Charles Darwin recognized that cooperation posed a challenge to his theory of natural selection. He described an elegant experiment to ferret out whether the aphid yields its excretion to the ant voluntarily, or involuntarily with the ant as a parasite.<sup>6</sup> He provided a convincing argument that it was the former. Darwin, the consummate naturalist, hated overgeneralized theory. Yet the significant literature on the evolution of cooperation that has flourished in the last fifty years is almost entirely theoretical. Much of it is cast in the formalism of mathematical game theory, a subject that came into existence more than half a century after Darwin’s death in the work of John von Neumann and Oskar Morgenstern. Game theory describes how competing, sentient players, in a well-de-

finer universe of choices and payoffs, may knowingly seek to optimize their own outcomes. Evolution is the blind watchmaker,<sup>7</sup> optimizing only by trial and error. Exactly how the achievable outcomes of evolution correspond to the mathematical optima of game theory is not a settled question.

**G**O BACK TO MICROBES *A* and *B*, but now promote them to sentience. They become Alice and Bob, who are arrested on suspicion of committing, together, a serious crime. Each has sworn not to betray the other. They are questioned in separate rooms.

“We already have enough evidence to convict you both of a misdemeanor,” the detective says to each, “that will put you away for one year.” Each, separately, says nothing. “But if you defect, rat out your partner and turn state’s evidence,” the detective continues, “we’ll let you go, scot-free. Your partner will get a felony conviction, six years in the state penitentiary.”

“What if we *both* turn state’s evidence?” Alice and Bob each ask.

“Well, I can’t let you both go free,” the detective says. “You’ll each get three years.”

Alice reasons as follows: there are only two possibilities. Either Bob will rat me out, or else he won’t. If he rats, then I’ll get six years—unless I rat also, in which case I’ll get just three years. So, if he rats, I should too. But what if he doesn’t rat? What a chump! I can rat on him and be out today. So, *either way*, I should defect. Bob employs the same reasoning and defects on Alice. They each get three years. The pair spend the time wishing that they had both kept their promises not to betray each other and escaped with misdemeanor convictions.

The prisoner’s dilemma (PD) game, played once, has no direct bearing on evolution. But consider the *iterated* prisoner’s dilemma (IPD) game, first posited at the RAND Corporation in the 1950s: Alice and Bob play many rounds of the same game with each other. After following the same logic for a few games, Alice tentatively tries a round of cooperation. In that round, Bob still defects and Alice gets six years. But Bob has now seen Alice’s signal. He tries cooperation himself. Alice reciprocates. And, for a string of games, they are both cooperating, receiving only misdemeanor convictions. In the IPD, there is information in the previous plays, and each player can try to use that information to devise a superior strategy that remains self-interested.

Is the best strategy to cooperate always? Certainly not. If Alice adopts that strategy, Bob will always defect, getting off scot-free, while Alice will always get six years. A good strategy would seem to be something like, “Cooperate most of the time, but don’t be a chump if the other player doesn’t follow suit.” Can this be formalized, or made crisp, in some way? By one estimate, more than 200 experiments on IPD, with human or computer players, were

conducted between 1965 and 1971. Robert Axelrod called IPD “the *E. coli* of social psychology.”<sup>8</sup> Axelrod’s own 1980 experiments are the most famous.<sup>9</sup> Human contestants submitted algorithmic strategies that, programmed on a computer, were all played against each other in a tournament. A strategy could be very complex. Alice could, in principle, look at Bob’s previous thousand plays and try to crack the code of his decision process. It was a big surprise when a very simple strategy, known as tit for tat (TFT), won the Axelrod tournaments. TFT starts by cooperating. Then, if the other player cooperates, it cooperates on the next round. If the other player defects, then, on the next round, it defects.

In later experiments, a related strategy, generous tit for tat (GTFT), was found to do even better. GTFT is the same as TFT, except that, when a tit-for-tat defection would be prompted, GTFT *sometimes*, and with a fixed probability, cooperates rather than defects—it is generous in that way. There seemed to be moral lessons in these results, pseudo-mathematical justifications of high human values. TFT embodied the Golden Rule. GTFT went further: turn the other cheek.

Science and ethics were in harmony.

I had long ago read Axelrod’s book and William Poundstone’s popular exposition, and I knew about TFT and GTFT.<sup>10</sup> That was about all I knew. During the slow period between Christmas and New Year’s in 2011, I set out to understand why TFT and GTFT did so well. I was struck by the fact that they both were memory-one strategies. That is, a player’s move—cooperate or defect—depended only on the immediately previous round of play. Had anyone ever proved that GTFT, or any other strategy for that matter, was optimal among memory-one strategies? If so, I couldn’t find it in Google Scholar.

Without getting into too much detail, every memory-one strategy is defined by four numbers, probabilities in the range between zero and one. A match between two players, each playing a fixed memory-one strategy, is thus described by eight parameters—equivalent to a point inside, or on the surface of, an eight-dimensional hypercube. It is not hard to derive formulas for the statistically expected gain or loss of each player as a function of the match position in the hypercube. With these formulas, it is not necessary to play out the games at every point, a huge saving in computer time.

Game theory has the important concept of Nash equilibrium. As the players each try to improve their strategies, one imagines the match-strategy point moving around in the hypercube. A Nash equilibrium is a point where no small adjustment of Bob’s four parameters by Bob improves his score, and the same is true for Alice’s four parameters and her score. If the players reach a Nash equilibrium, then neither will see any benefit in changing strategy further. Each sees an optimum that is at least local. Over the holiday, I wrote what I thought was an elegant computer

program to explore the game hypercube, always seeking better strategies. I expected it to find that the Nash equilibrium was something like GTFT for both players.

On the computer screen, I watched the progress of the optimization. The match seemed to be getting closer and closer to a Nash equilibrium—until the program crashed. I tried again, starting from a different point in the hypercube. Again, the program crashed, but at a different place. I automated the procedure and ran the program a thousand times—and got a thousand crashes. But they weren't at random points in the eight-dimensional hypercube. All of the crashes occurred on a particular four-dimensional hyperplane. I traced their cause to a faulty assumption in my program. It assumed that when Bob changed his strategy, it would, generically, have some effect on his own score, and similarly for Alice. How could this not be true? Apparently it wasn't. The computer could find instances, but not explain them.

That was when I knew that I needed Freeman Dyson. The exact complement to computer intelligence, as yin to yang, is Freeman Dyson intelligence. This problem needed both kinds. I emailed a description of my puzzling results to Freeman. A day later, he sent back a note with the general result all worked out, which became equations 1 through 7 in our paper.<sup>11</sup> When the confusion was hacked away, it came down to a simple equation in high-school algebra, as Freeman noted. Hidden within the IPD was a matrix whose determinant could be forced, by either Alice or Bob acting alone, to be zero—with quite unexpected implications. We called the resulting strategies zero determinant (ZD).

Overlooked in fifty years of research on IPD was the simple fact that Alice had the ability, by choosing a certain ZD strategy, to set Bob's score. She could pick any value between that of full cooperation (the one-year misdemeanor) and full defection (the three-year felony). No matter how Bob played, that value, on average, would be his outcome. And Bob, correspondingly, could do the same thing to Alice. Game theorists already had a name for this situation: an ultimatum game. They had no idea that there was an ultimatum game hidden inside the IPD.

In the classic ultimatum game, a hundred dollars appears out of nowhere on the table between Alice and Bob. By a coin flip, one player, say Alice, goes first. "I'll take \$60, and you can have \$40," she tells Bob. He can either accept and take the \$40, or else the whole hundred disappears and neither gets anything. What makes it a game is that Alice can pick any number between \$0 and \$100 for her initial offer. What is her optimal strategy? Mathematical game theory doesn't provide an answer; it turns out that all values are Nash equilibria, but not in a useful way. Psychologists and economists have tried the ultimatum game experimentally across a wide range of locations and ethnic groups.<sup>12</sup> The first player never offers more than \$50—why should she?—and the second player rarely accepts less

than \$20 or \$30, so a 60:40 split is quite typical. The game has been played between humans, and chimpanzees using raisins, with ambiguous results.<sup>13</sup>

The ultimatum variant revealed by the ZD strategies could be this: "Set my score to the lightest-sentence misdemeanor," Alice tells Bob, "and I'll do the same for you. But if for any reason you cross me, I'll change my strategy to punish you severely." This is reminiscent of TFT, but it is played at the meta-level of altering strategies, not at the game level of individual moves. To see the difference, imagine a scenario in which Alice is facing Darwin rather than Bob. Darwin is the blind watchmaker who can only try, by small mutations in strategy, to maximize his score. He doesn't know that he is playing against a sentient being. Evolutionary biologists speak of the fitness landscape of hills and valleys in which evolution by natural selection can be viewed as taking place. But here, Alice completely controls Darwin's fitness score. She can simulate any fictitious evolutionary landscape she wants. Darwin cannot distinguish it from a natural one. Thanks to ZD, she completely controls the apparent rules of the game. Biologists like to say, "You can't fool evolution," but this example shows that, within the constructed space of IPD, the ZD strategies actually can. That surprised a lot of people.<sup>14</sup>

The ZD strategies had yet more surprises in store. Instead of setting Bob's score to a value, there is a ZD strategy by which Alice can set Bob's score to be related to her own, extortionately. Alice picks an extortion factor: five, for example. She implements the corresponding ZD strategy for her own game, then lets Bob play any strategy he wants. If his strategy reduces his felony sentence by an amount  $x$ , which we will call his bonus, then Alice's bonus will be  $5x$ . If he optimizes his strategy by blind evolution, then Alice will score better than even if both cooperated—not even a misdemeanor sentence, but only a traffic violation. Bob ends up worse off. There is nothing he can do about it. Alice sets the strategy.

Really? Nothing? Can't he do back to Alice exactly the same thing as she is doing to him? Yes, he can. He can adopt a ZD strategy that sets *his* bonus to be five times *hers*. How can each of them have five times the bonus of the other? Easy: both get zero. In effect, a return to noncooperating three-year felony sentences. So, they are back to bargaining, as in the ultimatum game: "If I do this for you, will you do this for me?"

The most interesting result Dyson and I reached was this: that the outcome of the game depended on whether each player had a theory of mind about the other. Psychologists use the term *theory of mind* to mean the ability to attribute mental states—such as belief and intention—to others. When Alice plays against Darwin, Darwin has no theory of mind. Alice can extort him, or lead him around a fictitious fitness landscape, at will. But when Alice plays against Bob, he may have a theory of mind. Noticing that

his opponent's bonuses are always five times his own, he thinks, "I am being extorted! I will make us both get zero. *She will notice this.* Then, we can do a deal." But Bob may be wrong about this. Alice may have set her ZD strategy once and for all, put it on autopilot, and then vanished. Bob, with or without a theory of mind, can try anything he wants. He can only hurt himself. His best option is to accede to the extortion. "Press and I have solved the Prisoner's Dilemma game," Dyson told people jocularly. "The winning strategy is to *go to lunch.*"

**T**HE RESULTS DESCRIBED are so counterintuitive that some readers may want to get a sense of how they are derived. Label the four outcomes of the previous round 1, ..., 4 in the order (*cc, cd, dc, dd*), where *c* denotes cooperate, *d* denotes defect, and the order is Alice–Bob. Alice's strategy can be written as  $\mathbf{p} = (p_1, p_2, p_3, p_4)$ , her probabilities for cooperating on the current round given each of the four previous outcomes. Similarly, Bob's strategy is  $\mathbf{q} = (q_1, q_2, q_3, q_4)$ .

The strategies  $\mathbf{p}$  and  $\mathbf{q}$  imply a Markov process that advances any mixture of outcomes one round at a time. The Markov transition matrix is

$$\mathbf{M} = \begin{bmatrix} p_1q_1 & p_1(1-q_1) & (1-p_1)q_1 & (1-p_1)(1-q_1) \\ p_2q_2 & p_2(1-q_2) & (1-p_2)q_2 & (1-p_2)(1-q_2) \\ p_3q_3 & p_3(1-q_3) & (1-p_3)q_3 & (1-p_3)(1-q_3) \\ p_4q_4 & p_4(1-q_4) & (1-p_4)q_4 & (1-p_4)(1-q_4) \end{bmatrix}.$$

Because every Markov matrix has a unit eigenvalue, the matrix  $M' \equiv M - I$  is singular and has zero determinant. The stationary vector, i.e., the long-term probability mix of outcomes of the game, satisfies

$$\mathbf{v}^T \mathbf{M} = \mathbf{v}^T, \text{ or } \mathbf{v}^T \mathbf{M}' = 0.$$

The significance of the stationary outcome probability vector is that its dot product with Alice's prison sentences  $\mathbf{SA} = (1, 0, 6, 3)$  or Bob's  $\mathbf{SB} = (1, 6, 0, 3)$  gives the expectation value of their respective times in jail. Close attention should be paid to the extent that each may be able to unilaterally influence  $\mathbf{v} \cdot \mathbf{f}$  for a given  $\mathbf{f}$ , where  $\mathbf{v}$  is the stationary vector and  $\mathbf{f}$  is a given four-vector.

Now the promised high-school algebra: Cramer's rule for calculating a determinant, applied to the matrix  $M'$ , is

$$\text{Adj}(M')M' = \det(M')I = 0,$$

where  $\text{Adj}(M')$  is the adjugate matrix—what most of us in high school learned as the matrix of minors. This equation implies that every row of  $\text{Adj}(M')$  is proportional to  $\mathbf{v}$ . Choosing the fourth row, we see that the components of  $\mathbf{v}$  are, up to a sign, the determinants of the  $3 \times 3$  matrices formed from the first three columns of  $M'$ , leaving out each one of the four rows in turn. These determinants

are unchanged if the first column of  $M'$  is added into the second and third columns.

The result of these manipulations is a formula for the dot product of an arbitrary four-vector  $\mathbf{f}$  with the stationary vector  $\mathbf{v}$  as a single  $4 \times 4$  determinant:

$$\mathbf{v} \cdot \mathbf{f} = \det \begin{vmatrix} -1 + p_1q_1 & -1 + p_1 & -1 + q_1 & f_1 \\ p_2q_3 & -1 + p_2 & q_3 & f_2 \\ p_3q_2 & p_3 & -1 + q_2 & f_3 \\ p_4q_4 & p_4 & q_4 & f_4 \end{vmatrix}.$$

It is now possible to see the remarkable fact that the second column is entirely under Alice's control (the  $\mathbf{p}$ 's), while the third column is under Bob's control. Given a particular  $\mathbf{f}$ , each has the possibility of choosing a strategy that will make the determinant zero.

From here, it is only a short step to the results already mentioned. Alice chooses for  $\mathbf{f}$  a desired linear combination of the score vectors,  $\mathbf{f} = \alpha \mathbf{SA} + \beta \mathbf{SB} + \gamma$ , and then calculates a strategy  $\mathbf{p}$  that zeros the determinant. That linear combination of scores is made zero. These are exactly the ZD strategies, including the specializations of extortionate or ultimatum.

Our paper was published in *Proceedings of the National Academy of Sciences* in mid-2012. Despite a simultaneously published, glowing commentary by evolutionary biologists Alex Stewart and Josh Plotkin, few people seemed to care about the theory of mind.<sup>15</sup> What attracted attention was the extortion business. People reacted to it emotionally. Poundstone posited links to abusive marriages, terrorism, the current US Congressional deadlock, and income inequality.<sup>16</sup> Commenting on an article about the paper, a reader remarked, "This and similar 'quant' nonsense is precisely what has led to the ... too-big-to-fail banking disaster we are currently confronting. ... These studies are worse than useless, they are parasitic cancers on society."<sup>17</sup> Elsewhere, a kinder commenter added, "Once again, physicists invade a field and add value." Even that seemed barbed.

All this from a simple equation of high-school algebra.

Emotion tainted even the soberest responses to our work. Mathematically speaking, we had shown that a broad class of two-person games—not just IPD—had an unsuspected hidden algebraic structure, our ZD. It turned out that there were other, quite different, hidden structures. Mathematician Ethan Akin discovered a set of strategies with the twin properties that Alice and Bob both got off with minimal misdemeanors, and that neither could gain by unilaterally changing strategy.<sup>18</sup> He named these *good* strategies. Some of our ZD strategies were *good*, but others were not. By implication, they were evil. Stewart and Plotkin wrote an elegant paper focusing on the subset of ZD strategies that were *generous*, meaning that one player voluntarily ceded a greater share of reward to the other.<sup>19</sup> I was relieved that some of our ZD strate-

gies turned out to be both good and generous, though not including, of course, our extortionate strategies—which, by the way, would still beat all the good and generous strategies in head-to-head competition.

Alice and Bob play against each other over and over, but that is not how populations work. Population biologists study the dynamics and evolution of a mix of strategies, where an Alice must sometimes compete against a Bob, or another Alice, or the mutant Alicia—like Alice, but subtly different. Additional concepts come into play. Is the mutant able to *invade* the population, so that, by its greater fitness over many generations, its Alicias become dominant? The Bobs may have a strategy that is successful when played against the Alicies, but so mutually destructive when played against each other that a population dominated by Bobs is not an *evolutionary stable strategy*. In a finite population, there is a related kind of evolutionary inertia: a favorable mutation must exceed a certain threshold to avoid, on average, dying out just by chance. Bob's, or Alice's, strategy may win but still not be evolutionarily robust.

To the seeming delight of many, the extortionate ZD strategies were found to be neither stable nor robust, nor able to invade a population to any significant degree. In a population, extortion was, roughly speaking, self-limiting. Extortioners would tend to mutate into generous players because, most of the time, generous players would be playing against other generous players. Christoph Adami and Arend Hintze later quipped that this proved mathematically that winning isn't everything.<sup>20</sup> Dyson liked these results. I thought they were interesting and advanced the field, but I was bothered by the emotional coloration that seemed to accompany them. Freeman, I concluded, had gone over to the side of the sentimentalists.

We can ponder whether anything of value can be learned from the events I have described. That a couple of physicists—the more senior, eighty-eight years old—could invade a field and, over a holiday vacation, find an undiscovered nugget capable of attracting such attention may say something about serendipity, or about the genius of Freeman Dyson; or it may suggest that subfields of science can easily become too set in their ways, and that the scientific enterprise should seek institutional mechanisms that encourage more cross-fertilization over scientific boundaries.

Also worth pondering is the human tendency to label scientific findings with emotive words like good, generous, and, yes, extortionate. Most of the time this surely does no harm. It makes the science livelier and helps communication to each other and the public. Occasionally, though, the labels in a field become a mythos that can color it with subjectivity. The application of game theory to evolutionary biology led to a mythos, according to Poundstone, that you can't fool evolution and the most successful strategies are fair ones.<sup>21</sup> Neither assertion is a scientific truth. In the nineteenth century, one particular mythos attached

to natural selection was that it wasn't just a description of the way things *were*, but a description of the way things *ought to be*. That led to social Darwinism and its misuse of science in justification of racism, imperialism, eugenics, and other horrors. Evolution by natural selection *is what it is*. Everyone should be on guard against labeling it with either moral virtues or moral failings. A decade before the publication of Darwin's *On the Origin of Species*, Alfred Lord Tennyson's poem "In Memoriam A. H. H." gave a startlingly good description of natural selection. This is the poem where "Nature, red in tooth and claw" ultimately loses out to "...love, Creation's final law." But it is poetry, not science.

If I ever find myself fighting for survival in a population of Alicies, Bobs, Darwins, and Alicias, I plan to go with ZD extortion. And I would advise you to do the same—but just not too many of you.

*William Press is the Leslie Surginer Professor of Computer Science and Integrative Biology at the University of Texas at Austin.*



1. Carl Sagan, *Contact* (New York: Simon and Schuster, 1985).
2. Cixin Liu, *The Three-Body Problem*, trans. Ken Liu (New York: Tor Books, 2014).
3. An elegant review by Martin Nowak ("[Five Rules for the Evolution of Cooperation](#)," *Science* 314, no. 5,805 (2006): 1,560–63, doi:10.1126/science.1133755) compares five mechanisms: direct reciprocity, indirect reciprocity, kin selection, network reciprocity, and group selection.
4. Richard Dawkins, *The Selfish Gene* (New York: Oxford University Press, 1976).
5. Nowak, "[Five Rules](#)," 1,560.
6. Charles Darwin, Chapter VIII, in *On the Origin of Species* (London: John Murray, 1872 [1859]).
7. The blind watchmaker metaphor was introduced by William Paley in his 1802 book *Natural Theology* (London: R. Faulder), long before Darwin.
8. William Poundstone, *Prisoner's Dilemma* (New York: Anchor/Doubleday, 1992).
9. Robert Axelrod, *The Evolution of Cooperation* (New York: Basic Books, 1984).
10. Axelrod, *Evolution of Cooperation*. Poundstone, *Prisoner's Dilemma*.
11. William Press and Freeman Dyson, "[Iterated Prisoner's Dilemma Contains Strategies that Dominate Any Evolutionary Opponent](#)," *Proceedings of the National Academy of Sciences* 109, no. 26 (2012): 10,409–13, doi:10.1073/pnas.1206569109.
12. Hessel Oosterbeek, Randolph Sloof, and Gijs van de Kuilen, "[Cultural Differences in Ultimatum Game Experiments](#):"

- [Evidence From a Meta-Analysis](#),” *Experimental Economics* 7, (2004): 171–88, doi:10.1023/B:EXEC.0000026978.14316.74.
13. Joseph Henrich and Joan Silk, “[Interpretative Problems with Chimpanzee Ultimatum Game](#),” *Proceedings of the National Academy of Sciences* 110, no. 33 (2013): E3049, doi:10.1073/pnas.1307007110.
  14. William Poundstone, “[On ‘Iterated Prisoner’s Dilemma Contains Strategies that Dominate Any Evolutionary Opponent’ by William Press and Freeman Dyson](#),” *Edge.org*, June 18, 2012.
  15. Alexander Stewart and Joshua Plotkin, “[From Extortion to Generosity, Evolution in the Iterated Prisoner’s Dilemma](#),” *Proceedings of the National Academy of Sciences* 110, no. 38 (2013): 15,348–53, doi:10.1073/pnas.1306246110.
  16. Poundstone, “[On ‘Iterated Prisoner’s Dilemma’](#).”
  17. Although the original article posted by *MIT Technology Review* in 2012 is no longer available, this response can be found in [an archived discussion thread](#) on the website *Hacker News*, dated August 17, 2012. See [the comment posted by the user ramblerman](#).
  18. Ethan Akin, “The Iterated Prisoner’s Dilemma: Good Strategies and Their Dynamics,” arXiv, arXiv:1211.0969v3.
  19. Stewart and Plotkin, “[From Extortion to Generosity](#).”
  20. Christoph Adami and Arend Hintze, “[Evolutionary Instability of Zero-Determinant Strategies Demonstrates That Winning Is Not Everything](#),” *Nature Communications* 4, no. 2,193 (2013), doi:10.1038/ncomms3193.
  21. Poundstone, “[On ‘Iterated Prisoner’s Dilemma’](#).”
- DOI: 10.37282/991819.21.1