How Real People Think in Strategic Games

By Sara Robinson

Once upon a time, Arianna, a Hollywood hostess, threw a large party for a diverse assemblage of guests. Bored with the usual rounds of "Hollywood Star Charades" and "Two Minutes with a Millionaire," she initiated a novel game: Each guest was to choose a number from 0 to 100, write it on a piece of paper, and drop it in a passed basket. The goal, she explained, was to choose a number as close as possible to two thirds of the group average. The person or people who came closest would share a bottle of very expensive champagne.

Minerva, the lone mathematician at the party, mulled over possible strategies. "Suppose the other participants' choices are randomly distributed," she reasoned. "The average of the other guesses would be about 50, so I should choose 33. . . . But wait, the other participants will realize this too, so their numbers would not be random—they'd average to 33. So I should choose 22. . . . But the others will see this too, so I should choose 14. . . ." As she quickly realized, this iterative process stops only when it reaches 0. In the end, assuming that the other guests would follow a similar line of reasoning, Minerva decided that the group average would be 0. Anticipating at least a taste of superb champagne, she confidently wrote down her guess and dropped it in the basket.

Minerva had happened on the unique "Nash equilibrium" for Arianna's game, although, as a number theorist, she wasn't familiar with the term. A Nash equilibrium is a collection of strategies, one for each player, such that even if all players know the others' strategies, they have no incentive to change their own.

A little later, Arianna tabulated the guesses and, with fanfare, announced the group average: 30. Minerva was stunned. Warren, the lone economist in the group (who had, in fact, suggested the game to the hostess), won the champagne by guessing 31. He said his farewells and left the party with the champagne and a well-known actress on his arm.

What was the error in Minerva's reasoning?

Game Theory for Real People

Minerva's mistake was her assumption that the other party guests would reason the way she did. While a few did guess very low numbers, many made random guesses between 0 and 100. Others started along Minerva's chain of thinking but stopped at 33, or 21. In fact, Minerva was the only one to guess 0.

In classical game theory, the assumption is that participants in a game will, like Minerva, reason rationally and thus play equilibrium strategies. But, as Arianna's game illustrates, people often don't behave as mathematical theory predicts.

"People have models of other people that are different than equilibrium models," says Vincent Crawford, a professor of economics at the University of California, San Diego. "Therefore, they may not play their equilibrium strategy, even if they understand the reasoning that leads to it."

Crawford is working in a relatively new area of game theory in which the aim is to better predict the way people actually think about strategic games. Known as behavioral game theory, this research area combines elements of psychology and neuroscience with traditional economic theory, emphasizing experimental studies as the best way to describe human behavior. (See James Case's "Economics as Lab Science: Nobel Prizes Give Boost to Growing Field," *SIAM News*, April 2003, http://www.siam.org/siamnews/04-03/tocapr03.htm.)

The experiments follow a pattern: The researchers use simple games for which classical game theory makes clear predictions about participants' behavior. In playing the games, subjects make simple choices, knowing that their choices will combine with those of other participants to produce a payoff. The subjects are then actually rewarded (with cash) according to their performance, and the experiments are repeated to see how the subjects learn over time. Using the results of these experiments, behavioral game theorists hope to develop new, more nuanced theories that are useful for applications.

One focus of behavioral game theorists has been "dominance solvable games"—games that, like Arianna's, can be solved by iterated reasoning. For mathematicians, the inductive leap from one or two iterations to the mathematical limit is an effortless one. But this is not the case for most people. From experimental studies, behavioral game theorists conclude that people tend to do only a few steps of iterated reasoning and then stop—either because the reasoning is too complicated for them or because they believe it's too complicated for others.

Beauty Contests and Iterated Reasoning

Arianna's party game is known to economists as "the beauty contest game," after an observation of John Maynard Keynes. In his 1936 book *General Theory of Employment, Interest, and Money*, Keynes compared the stock market to a newspaper beauty contest in which participants guess which of the photographed faces others will judge to be the most beautiful.

"It is not a case of choosing those which, to the best of one's judgment, are really the prettiest, nor even those which average opinion genuinely thinks the prettiest," Keynes wrote. "We have reached the third degree, where we devote our intelligences to anticipating what average opinion expects the average opinion to be. And there are some, I believe, who practice the fourth, fifth, and higher degrees."

First studied experimentally by Rosemarie Nagel, the beauty contest game appeals to behavioral economists because it is simple,

yet captures the reasoning described by Keynes, which is present in many real-world situations.

In 1995, Nagel had groups of German students, about fifteen in each group, try to guess two thirds of the group average within limits of 0 and 100. While the average guess for the groups was around 35, she observed that many students chose 33, two thirds of the midpoint of 50, or 22, two thirds of two thirds of the midpoint. Very few players chose 0.

Three years later, Teck Ho, Colin Camerer, and Keith Wiegelt replicated the experiment, replacing two thirds with 0.7, 0.9, 1.1, and 1.3. With their data, the researchers precisely quantified the number of steps of iterated reasoning the subjects seemed to be doing. Most subjects, they concluded, were doing between 0 and 3 rounds.

In a fascinating series of unpublished experiments, Camerer, a professor of economics at the California Institute of Technology, also tried the game with several distinct subject pools. When the goal was to guess 70% of the average, some of the groups, including those made up of Caltech undergraduates, game theorists, and computer scientists, averaged under 20, while other groups, such as CEOs, 70-year-olds, and Pasadena City College students, averaged over 50.

Again using a statistical analysis of the guesses, Camerer inferred the average number of "steps of thinking" for each of the groups; the numbers range from 0 (for a small group of Pasadena City College students) to 3.8 (for the computer scientists).

A Piece of the Pie

So-called "bargaining games" also illus-trate limited iterated reasoning, together with other features of behavioral game theory. In one simple example, two players are asked to apportion a $10 "pie." The first player is given the $10 and told to offer some fraction of it to the second player. The second player accepts or rejects the offer; in the latter case, he gets to make a counteroffer of a fraction of a $5 pie to the first player. If that offer is rejected, the first player gets to make a final apportionment of a $2.50 pie; if that is rejected, neither player gets anything.

A game theorist (or a mathematician) might reason through this game as follows: Suppose the game makes it to the last round. Because by then the second player has no leverage, the first player can offer only 1 cent and then keep $2.49. Thus, in round two, the second player should offer $2.50 to the first, to avoid going to round three. If he offers $2.50, he also keeps $2.50; in the first round, the first player should thus offer $2.51, getting to keep $7.49, to avoid going to round two.

In study after study, however, first-round offers tended to be far larger than $2.51, typically between $4 and $5, and for good reason: About half of offers of $2 or less are summarily rejected by the second player. Apparently, making money is not the players' only concern; participants have a sense of pride and care about how they are treated by others, economists have concluded. Thus, offers perceived to be "unfair" are rejected out of a desire for revenge.

Experiments show that limited iterated reasoning plays a role as well. Camerer, with Eric Johnson, Talia Rymon, and Sankar Sen, devised a study utilizing a computer interface to monitor subjects' searches for hidden but freely accessible payoff information. The interface enabled the researchers to track the pieces of information the subjects used and the order in which they used them. From this data, the researchers concluded that most players look ahead only one round. When players are trained to look further ahead, the researchers showed, they do make offers closer to the rational theory prediction.

For the most part, the behavior of participants is consistent across different cultures, even for large sums of money. In dozens of experiments conducted in such countries as Indonesia and Slovenia, subjects rejected low offers, even when the pie was as large as $400. Curiously, the only exceptions observed were in so-called "primitive" cultures, such as tribes in Africa, the Amazon, and Papua New Guinea. In these groups, subjects tended to offer very little, typically less than 15%, and responders accepted virtually every offer. Amazingly, people in these cultures seem to be unique in behaving exactly as game theory predicts.

Dominance Solvable Games: A Recent Experiment

An experimental as yet unpublished study by Miguel Costa-Gomes of the University of York and Vincent Crawford gives a taste of current research in the area of limited iterated reasoning.

Costa-Gomes and Crawford carefully designed a series of 16 two-player games, loosely modeled on the beauty contest game, but with a twist: Rather than asking all the players to get as close as possible, within a common pair of limits, to a common number p times the group average, the researchers assign each player in a game a number (called the player's target) and a pair of limits. A player's object is to get as close as possible, within her limits, to the product of her target and the other person's guess. Players don't have to guess within their allowed ranges, because guesses falling outside the limits are automatically adjusted to the limits. In the Costa-Gomes–Crawford games, the lower limits are always 100 or 300, and the upper limits 500 or 900. The targets are taken from the set {0.5, 0.7, 1.3, 1.5}.

Readers will quickly see that each game will have a unique equilibrium, determined (not necessarily directly) by the players' lower limits when the product of the targets is less than one and by the upper limits when the product is greater than one.

After reading the instructions and playing four practice rounds, the subjects play the 16 games in succession on a computer, receiving no feedback on one game before launching into the next. Thus, while participants are paired for the purpose of determining their payoffs, they get no feedback from their partners after the training session.

The way players think about the game is monitored with an interface similar to that of Camerer et al. in the bargaining experiment. The six pieces of information needed to play one game—the player's limits and target plus her opponent's limits and target—are given in an array of boxes. The number in each box is visible only when the player clicks on that box, and only one number can be viewed at a time. Participants are not allowed to use pencil and paper and so must look up each number as they need it. The sequence and duration of lookups are recorded by the computer.

From this data, the researchers concluded that subjects did from 0 to 3 steps of iterated thinking, but almost never chose equilibrium guesses. The data also revealed large "spikes" in the strategy spaces, indicating a nonrandom structure in the reasoning

strategies. Each spike, the researchers determined, corresponded to a pattern of decision rules for playing the game. The researchers associated each set of rules with an ideal "type," adapted from the literature on strategic decision-making. (The researchers' data, the instructions, and other details can be found at http://weber.ucsd.edu/~vcrawfor/#Guess.)

An L_1 player, for instance, following the most naïve strategy, assumes that her opponent's guesses are uniformly distributed over the allowed range and gives a best response to this scenario. An L_2 person assumes her opponent to be L_1 and gives a best response; an L_3 person assumes an L_2 opponent, and so on. An "equilibrium" person always gives the theoretically best equilibrium strategy, while a "sophisticated" player acts like a behavioral game theorist, making the guess that yields the highest possible expected payoffs, given others' guesses in the experiment (a proxy for the predictions of not-yet-developed behavioral game theory). The parameters of the games were chosen so as to maximally separate the types by their predicted guesses and lookup patterns.

A player's type can be inferred from her lookup sequence. A perfect L_1 player, for instance, takes the midpoint between her partner's limits and multiplies it by her target to obtain her guess. She doesn't use the partner's target in her calculation, nor does she need to look at her own limits—the program automatically adjusts the guess. The expected lookup sequence of an L_1 player would then be each of her partner's limits, followed by her own target.

A key advantage of this expanded form of the beauty contest game, Crawford explains, is its creation of a continuous, large strategy space that allows for separation of the different types' guesses. Indeed, a surprisingly large number of the subjects were true to a type, using exactly the lookup patterns expected for their types in 9 or more of the 16 games.

Players of L_k strategies mimic equilibrium strategies in simple games but deviate systematically from the equilibrium in more complex games, Crawford points out. He believes that an understanding of the structure of these deviations can enhance the predictive power of economic theory.

"Rules like L_k are simple heuristics for predicting others' strategic responses, while avoiding much of the cognitive complexity of equilibrium," Crawford says.

The Limits of Mathematics

While the mathematical models of von Neumann, Morgenstern, and others are aesthetically pleasing, and create a solid foundation for thinking about games, as these experiments and others show, elegant models often deviate from a more complex reality.

"Game theorists fell in love with the elegance of their model; it spoke to them and they thought that's how people really behave," Camerer says. "Over time, we've started to move away from a hyper-rational model to a more psychologically nuanced model."

Sara Robinson is a freelance writer based in Pasadena, California.