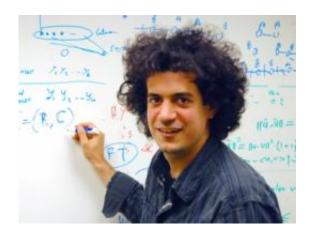# What computer science can teach economics

November 9 2009, by Larry Hardesty

Constantinos Daskalakis, an assistant professor in MIT's Computer Science and Artificial Intelligence Laboratory. Image: Satyen Kale

(PhysOrg.com) -- Computer scientists have spent decades developing techniques for answering a single question: How long does a given calculation take to perform? Constantinos Daskalakis, an assistant professor in MIT's Computer Science and Artificial Intelligence Laboratory, has exported those techniques to game theory, a branch of mathematics with applications in economics, traffic management -- on both the Internet and the interstate -- and biology, among other things. By showing that some common game-theoretical problems are so hard that they'd take the lifetime of the universe to solve, Daskalakis is suggesting that they can't accurately represent what happens in the real world.

Game theory is a way to mathematically describe strategic reasoning — of competitors in a market, or drivers on a highway or predators in a habitat. In the last five years alone, the Nobel Prize in economics has twice been awarded to game theorists for their analyses of multilateral treaty negotiations, price wars, public auctions and taxation strategies, among other topics.

In game theory, a "game" is any mathematical model that correlates different player strategies with different outcomes. One of the simplest examples is the penalty-kick game: In soccer, a penalty kick gives the offensive player a shot on goal with only the goalie defending. The goalie has so little reaction time that she has to guess which half of the goal to protect just as the ball is struck; the shooter tries to go the opposite way. In the game-theory version, the goalie always wins if both players pick the same half of the goal, and the shooter wins if they pick different halves. So each player has two strategies — go left or go right — and there are two outcomes — kicker wins or goalie wins.

It's probably obvious that the best strategy for both players is to randomly go left or right with equal probability; that way, both will win about half the time. And indeed, that pair of strategies is what's called the "Nash equilibrium" for the game. Named for John Nash — who taught at MIT and whose life was the basis for the movie A Beautiful Mind — the Nash equilibrium is the point in a game where the players have found strategies that none has the incentive to change unilaterally. In this case, for instance, neither player can improve her outcome by going one direction more often than the other.

Of course, most games are more complicated than the penalty-kick game, and their Nash equilibria are more difficult to calculate. But the reason the Nash equilibrium is associated with Nash's name — and not the names of other mathematicians who, over the preceding century, had described Nash equilibria for particular games — is that Nash was the

first to prove that every game must have a Nash equilibrium. Many economists assume that, while the Nash equilibrium for a particular market may be hard to find, once found, it will accurately describe the market's behavior.

Daskalakis's doctoral thesis — which won the Association for Computing Machinery's 2008 dissertation prize — casts doubts on that assumption. Daskalakis, working with Christos Papadimitriou of the University of California, Berkeley, and the University of Liverpool's Paul Goldberg, has shown that for some games, the Nash equilibrium is so hard to calculate that all the computers in the world couldn't find it in the lifetime of the universe. And in those cases, Daskalakis believes, human beings playing the game probably haven't found it either.

In the real world, competitors in a market or drivers on a highway don't (usually) calculate the Nash equilibria for their particular games and then adopt the resulting strategies. Rather, they tend to calculate the strategies that will maximize their own outcomes given the current state of play. But if one player shifts strategies, the other players will shift strategies in response, which will drive the first player to shift strategies again, and so on. This kind of feedback will eventually converge toward equilibrium: in the penalty-kick game, for example, if the goalie tries going in one direction more than half the time, the kicker can punish her by always going the opposite direction. But, Daskalakis argues, feedback won't find the equilibrium more rapidly than computers could calculate it.

The argument has some empirical support. Approximations of the Nash equilibrium for two-player poker have been calculated, and professional poker players tend to adhere to it — particularly if they've read any of the many books or articles on game theory's implications for poker. The Nash equilibrium for three-player poker, however, is intractably hard to calculate, and professional poker players don't seem to have found it.

How can we tell? Daskalakis's thesis showed that the Nash equilibrium belongs to a set of problems that is well studied in computer science: those whose solutions may be hard to find but are always relatively easy to verify. The canonical example of such a problem is the factoring of a large number: The solution seems to require trying out lots of different possibilities, but verifying an answer just requires multiplying a few numbers together. In the case of Nash equilibria, however, the solutions are much more complicated than a list of prime numbers. The Nash equilibrium for three-person Texas hold 'em, for instance, would consist of a huge set of strategies for any possible combination of players' cards, dealers' cards, and players' bets. Exhaustively characterizing a given player's set of strategies is complicated enough in itself, but to the extent that professional poker players' strategies in three-player games can be characterized, they don't appear to be in equilibrium.

Anyone who's into computer science — or who read "Explained: P vs. NP" on the MIT News web site last week — will recognize the set of problems whose solutions can be verified efficiently: It's the set that computer scientists call NP. Daskalakis proved that the Nash equilibrium belongs to a subset of NP consisting of hard problems with the property that a solution to one can be adapted to solve all the others. (The cognoscenti will infer that it's the set called NP-complete; but the fact that the Nash equilibrium always exists disqualifies it from NP-completeness. In fact, it belongs to a different set, called PPAD-complete.)

That result "is one of the biggest yet in the roughly 10-year-old field of algorithmic game theory," says Tim Roughgarden, an assistant professor of [computer science](#) at Stanford University. It "formalizes the suspicion that the Nash equilibrium is not likely to be an accurate predictor of rational behavior in all strategic environments."

Given the Nash equilibrium's unreliability, says Daskalakis, "there are

three routes that one can go. One is to say, We know that there exist games that are hard, but maybe most of them are not hard." In that case, Daskalakis says, "you can seek to identify classes of games that are easy, that are tractable."

The second route, Daskalakis says, is to find mathematical models other than Nash equilibria to characterize markets — models that describe transition states on the way to equilibrium, for example, or other types of equilibria that aren't so hard to calculate. Finally, he says, it may be that where the Nash equilibrium is hard to calculate, some approximation of it — where the players' strategies are almost the best responses to their opponents' strategies — might not be. In those cases, the approximate equilibrium could turn out to describe the behavior of real-world systems.

As for which of these three routes Daskalakis has chosen, "I'm pursuing all three," he says.

Provided by Massachusetts Institute of Technology ([news](#) : [web](#))