

The Work of Constantinos Daskalakis

Allyn Jackson

How long does it take to solve a problem on a computer? This seemingly innocuous question, unanswerable in general, lies at the heart of computational complexity theory. With deep roots in mathematics and logic, this area of research seeks to understand the boundary between which problems are efficiently solvable by computer and which are not.

Computational complexity theory is one of the most vibrant and inventive branches of computer science, and Constantinos Daskalakis stands out as one of its brightest young lights. His work has transformed understanding about the computational complexity of a host of fundamental problems in game theory, markets, auctions, and other economic structures. A wide-ranging curiosity, technical ingenuity, and deep theoretical insight have enabled Daskalakis to bring profound new perspectives on these problems.

His first outstanding achievement grew out of his doctoral dissertation and lies at the border of game theory and computer science. Game theory models interactions of rational agents—these could be individual people, businesses, governments, etc.—in situations of conflict, competition, or cooperation. A foundational concept in game theory is that of *Nash equilibrium*. A game reaches a Nash equilibrium if each agent adopts the best strategy possible, taking into account the strategies of the other agents, so that there is no incentive for any agent to change strategy. The concept is named after mathematician John F. Nash, who proved in 1950 that all games have Nash equilibria. This result had a profound effect on economics and earned Nash the Nobel Prize in Economic Sciences in 1994.

Game theory is widely applicable, not only in economics but even in such areas as international relations and biology. Clearly it would be useful to have a way to calculate Nash equilibria efficiently, in order to predict what strategic agents might do. Thus a long line of research began, soon after Nash proved his theorem, in which researchers developed algorithms to calculate Nash equilibria. By the early 2000s, however, none of those algorithms was known to be efficient, and some had been shown to be *inefficient*. Suspicions arose that the problem of computing a Nash equilibrium might be intractable. And if it is intractable, there would be no reason to expect equilibria would always be discovered by strategic agents, that is, by human beings with limited brains. Such a conclusion would reduce expectations for how well

Nash equilibrium can predict human behavior.

Questions of computational tractability are often viewed in the context of the well known “P versus NP” paradigm, which arose in the 1970s and soon became a central theme of computational complexity theory. Loosely speaking, P stands for the class of problems that are easy to solve by computer, meaning that an efficient algorithm for their solution is known. The class NP by contrast contains problems that are believed to be hard to solve, meaning that, if one is given a proposed solution, there is an efficient way to check it, but no efficient algorithm is known to produce solutions.

Problems in NP derive their hardness from the possibility that a solution might not exist. By contrast, for the problem of computing a Nash equilibrium, Nash’s proof guarantees that a solution exists. For this reason, the Nash equilibrium problem does not fit into the P versus NP paradigm. In 1994, Christos Papadimitriou defined a new complexity class called PPAD, suited to problems like the Nash equilibrium problem, for which solutions always exist and for which no efficient algorithm is known to compute solutions. PPAD stands for “polynomial parity argument for directed graphs” and refers to a certain standard argument used to prove existence results in combinatorics. The argument is a directed version of a result known as the handshaking lemma. PPAD contains all computational problems whose solution can be shown to exist by using this lemma.

One of the most important problems in PPAD is a computational version of a result from pure mathematics called the Brouwer fixed point theorem. It says that a continuous mapping from a ball to itself cannot displace all points; at least one point must remain fixed under the map. Proved by L.E.J. Brouwer in 1911, this fundamental result is the basis for countless proofs in mathematics—including Nash’s proof of the existence of equilibria. Brouwer’s proof is nonconstructive, meaning that it guarantees the existence of fixed points but does not show you how to find them. The computational version of the Brouwer fixed point theorem asks for an algorithm for finding fixed points. In his 1994 paper, Papadimitriou showed that this computational version is “PPAD-complete,” meaning that it is in PPAD and any problem in PPAD can be reduced to it, or, in other words, it is exactly as hard as PPAD.

Ten years later, Daskalakis became a PhD student of Papadimitriou and began working on the Nash equilibrium problem. He made a big advance when, together with Papadimitriou and Paul Goldberg, he proved that the Nash equilibrium problem is computationally equivalent to the problem of

finding Brouwer fixed points and therefore is also PPAD-complete. In showing that the Nash equilibrium problem is intractable, the work of Daskalakis et al breaks the universality of Nash equilibria: One cannot expect Nash equilibria to always result from interactions among strategic agents, because those agents cannot perform intractable computations. In this practical sense, Nash equilibria do *not* always exist.

The work also sheds light on why an efficient algorithm for the Nash equilibrium problem had been so elusive. If one looks at the guts of the algorithms people had developed to compute Nash equilibria, one sees the hallmark structure of the PPAD class lurking in the background. The work of Daskalakis et al showed that this structure is unavoidable.

While contributing new conceptual and technical insights, their work also emphasized the need for practical algorithms for efficiently computing equilibria for important subclasses of games. In later work, Daskalakis produced significant results on approximation of Nash equilibria, which have in turn provided inspiration for further developments by other researchers.

The second main area where Daskalakis has made outstanding contributions is connected to a branch of economics called mechanism design. Here a “mechanism” is a set of incentives that are offered in order to induce agents to behave in a certain way. Mechanism design is in some sense the reverse of game theory, in that game theory tries to analyze how agents will behave in a game, whereas mechanism design aims at designing a game that provides the right incentives for the agents to reach a desirable outcome.

The most basic model of a mechanism is an auction, and the simplest example is an auction with only one item for sale. How should one design the auction rules to maximize profit? In 1981 economist Roger Myerson provided a complete and elegant answer to this question for a single-item auction as well as for other auction settings in which each bidder’s preferences for the outcome can be summarized in a single number. Such auction settings are called “single-dimensional.”

Myerson’s work had a major impact in economics and stimulated a great deal of subsequent research. Its insights have been applied in a host of enterprises, such as auctions for drilling rights, telecommunications spectra auctions, and online auctions. The 2007 Nobel Prize in Economic Sciences honored the area of mechanism design and was awarded jointly to Myerson, Leonid Hurwicz, and Eric Maskin.

There are other kinds of auction settings, in which more than one item is offered for sale and the items might be bundled, and in which the value

the bidders put on the bundles cannot be expressed by a single number. In contrast to our thorough knowledge of single-dimensional auctions, such multi-dimensional auctions have remained poorly understood. In fact, it is not known how to optimally sell even two items to one buyer. As the number of buyers and items increases, the number of possible auction designs quickly proliferates into a highly complicated set. The optimal designs, when they can be found, often exhibit counter-intuitive properties and can be highly sensitive to details of the bidders' preferences.

Daskalakis delved into this dauntingly complex subject starting around 2011. One of the biggest challenges was to develop insight into the set of all possible auction designs. Daskalakis, together with his students at the time, Yang Cai and Matt Weinberg, developed ingenious ways to exploit linear programming to uncover structure in this set. Using this structure, they developed a method for translating a mechanism design problem into a problem in algorithm design. They could then construct computationally efficient algorithms to produce optimal mechanisms. The balance this work strikes between the structural and the computational viewpoints proved especially fruitful and yielded powerful results.

The new insights developed in this work led to further advances by Daskalakis, Alan Deckelbaum, and Christos Tzamos. They obtained new results in an area of mathematics known as optimal transport theory and used those results to characterize mathematically the structure of optimal multi-item mechanisms in the single-buyer setting.

The most valuable aspect of the work of Daskalakis in mechanism design has been to crack open a problem that had previously been seen as unapproachable. Because they are theoretical, the results would need to be simplified before they could be applied directly to concrete problems. Researchers, including Daskalakis, have already begun to explore ways to adapt his structural and algorithmic results to exchange optimality for simplicity and robustness.

In addition to his work on the Nash equilibrium problem and mechanism design, Daskalakis has made contributions to several other areas, including machine learning, statistics and probability theory, and computational biology. In the latter realm, a major task is the reconstruction of phylogenetic trees from molecular data. In 2011, together with Elchanan Mossel and Sebastien Roch, Daskalakis published a proof of a central conjecture in mathematical phylogeny concerning the specific conditions under which evolutionary trees can be reconstructed. His more recent work focuses on

high-dimensional statistics and the theoretical foundations of machine learning.

The work of Constantinos Daskalakis exhibits fearlessness in tackling difficult, complex, and longstanding problems. He engages deeply with their concrete details and uses the intuition thereby gained to synthesize structural and technical insights that provide the key to theoretical advances. His brilliance as a researcher combines with his lively curiosity and infectious enthusiasm to make him a natural leader in the field. Now just 37 years old, Daskalakis is sure to remain in this role in the decades to come.