Abstract

Traditional auctions, and more generally, traditional economic transactions, have always been of "manual scale." For example, the number of auctions conducted even in the largest of auction houses was moderate enough (and each item—expensive enough) so that expert appraisers could dedicate time to surveying each item sold, and use their expertise to determine a starting price that is optimized to maximize the revenue of the auction house. Such manual solutions were tangible until the Internet changed our world. For instance, whenever you search for something using your favorite search engine, a split-second auction is performed among advertisers who are related to your search term in order to determine which ads will be shown to you in the results page a mere moment later; the quantity of these auctions (billions per day), together with the low worth of each (many times less than a single cent each, yet in total reaching many billions of dollars), makes any manual intervention (say, in choosing the "starting price") completely impractical.

This explosion in online and computerized economic activity necessitates the study and understanding of economic mechanisms and markets of unprecedented scale. How good can simple mechanisms be? How complex must optimal mechanisms be? And, most substantially, what are the precise trade-offs between simplicity and quality? In this talk, I will survey some of my recent results on such questions for three notions of complexity within the context of the design of high-revenue auctions: the complexity of describing high-revenue auctions, the complexity of the machine-learning of high-revenue auctions, and the communication complexity of running high-revenue auctions.