Research Statement
Victor Naroditskiy

My main research interests lie at the intersection of computer science and economics. In particular, I am interested in computational approaches to mechanism design, sponsored search auctions, and automated bidding agents. In my research, I combine theory with practice. On the one hand, I approach theoretical problems computationally and use simulations to assist in the derivation of theoretical results. On the other hand, I design and implement autonomous agents to automate decision making in various trading domains.

Below I describe the main contributions of my Ph.D. research and future research goals. Broadly speaking, the contributions fall into two categories: computational approaches to economics problems and designing and implementing automated decision-making agents, particularly for trading domains.

Computational Approach to Economics Problems

Computational Approach to Designing Utility-Maximizing Mechanisms for Allocating Resources In the area of mechanism design, I apply a computational approach to obtain results that improve over those derived analytically. Solving for the best mechanism among those that meet relevant incentive constraints is a constrained optimization problem. Finding an analytical solution is most often too challenging, but some problems allow for a tractable algorithmic solution.

As an illustration of this line of research, consider a basic allocation setting where n agents commonly own $m < n$ identical items that they wish to distribute among themselves. Each agent wants exactly one item, and has a value for that item which is known to him alone. A mechanism is a means of distributing the items to the agents based on the values they report. Each agent is self-interested, so will misreport his value if doing so results in a higher individual utility. Payments can be introduced to make it in each agent's best interest to report his value truthfully. The agent's utility equals his private value minus the payment he makes if he gets an item. The agent's utility is just the payment he receives if he does not get an item. I am interested in designing a mechanism that maximizes some index of social welfare among those mechanisms that satisfy the relevant incentive constraints and guarantee voluntary participation.

A standard class of mechanisms for domains in which payments between agents are feasible is the Groves class. In this allocation domain, the items are commonly owned by the agents. Contrary to auctions, there is no auctioneer in this domain, and hence full efficiency is not compatible with incentive constraints - some money has to be burnt when implementing Groves mechanisms. Clearly burning payments is undesirable as it decreases total utility. Payments may be as high as the sum of the values of the agents who get the items; if these payments are burnt, total utility is zero.

In addition to Groves mechanisms that always distribute all items, my colleagues and I consider mechanisms that may choose to destroy some items in order to burn less money. The computational approach helps to explore the tradeoff between destroying items and burning payments. I show that destroying one item when $(m+1)$st highest value is much smaller than the mth highest results in a mechanism that burns much less money than the mechanism that does not destroy any items. I am able to significantly improve existing results when the number of items is not much smaller than the number of agents. Most notable are the results for the problems with $n-1$ items: the best Groves mechanism has the total utility of zero while the best mechanism I find is asymptotically optimal. For the case of 6 agents and 5 items, my mechanism achieves 93% of the maximum total utility, while the best known non-Groves mechanism achieves 36% and the best Groves mechanisms achieves 0%.

The work on mechanisms for allocating resources has a number of future directions. I am currently analyzing computational results in attempts to prove the optimality of the discovered mechanism. The next step is to investigate more general allocation settings characterized, for example, by the allocation of non-identical items, agents desiring more than one item, agents with utilities that depend on whether other agents receive items, and common-value models. Applying a computational approach to other domains where incentives are important is another avenue for future research. More generally, it is well understood what is implementable (e.g., Gibbard-Satterthwaite theorem, revelation principle). However less work has been devoted to finding the best (according to some metric) implementable mechanism. This is the question that computational techniques are well equipped to address.

Finding Bayesian Nash Equilibria in Auctions Computer science and economics are also interrelated in that theoretical results are evaluated in terms of the possibility of implementing them computationally. In particular, existence results may be supplemented by algorithmic solutions. Existence of Nash equilibria is well understood, but finding a Nash equilibrium is hard for a wide variety of non-trivial games. In particular, equilibrium strategies have been analytically derived only for a very small class of auctions. I tackle the problem of finding equilibria in auctions by investigating computational techniques that can help characterize equilibria. Specifically, I focus on symmetric Bayesian Nash equilibria in auctions and propose a computational method that takes advantage of the symmetry of equilibria and the structure of auctions. I present experimental results for single unit first- and second-price auctions with discrete values and bids.

Future research in this direction includes a computational search for equilibrium in simultaneous auctions for interdependent goods. Some of the methods that I will try are myopic best-response and variants of fictitious play.

Automated Decision Making

With my colleagues, I develop agents for automated decision making in three domains. Two of the domains are based on the annual Trading Agent Competition (TAC) and the third one is bidding in sponsored search auctions.

In the Travel division of TAC, an agent acts as a travel firm that provides travel packages for its clients. The problem the agent faces is what bids to place in a variety of simultaneous and sequential auctions for flight tickets, hotel reservations, and entertainment tickets. The goal is to satisfy clients' preferences while keeping the costs as low as possible. This problem is a model of many real-world problems, such as bidding for interdependent goods sold in different eBay auctions. I conduct a thorough investigation of different bidding strategies and design an agent that utilizes the Sample Average Approximation method to make bidding decisions under uncertainty about the auction clearing prices. Our agent won the final competition in 2006.

In the Supply Chain Management (SCM) division of TAC, an agent makes decisions for a computer manufacturing company that buys components, schedules assembly of computers, and sells finished products over the period of a simulated year. The main challenge in SCM is to coordinate long-term procurement and short-term selling decisions in the face of uncertainty about both future component costs and computer prices. I explore methods for addressing this challenge. The agents our team built were usually finalists in the SCM competitions.

The third problem domain that I research is that of sponsored search auctions. Sponsored search is a way for search engines to monetize search activity and an opportunity for advertisers to provide highly targeted ads. The ads are the sponsored links displayed on the search results page alongside the natural results for a search query. An advertiser needs to decide how to allocate its budget across different search queries. My co-author and I model this problem as a stochastic multiple-choice knapsack problem (S-MCKP) and propose a new algorithm to solve S-MCKP and the corresponding bidding optimization problem. The idea is to select items based on a threshold function which is generated and updated online as new sets of items arrive. The algorithm achieves about 99% performance compared to the offline optimum when applied to a real bidding data set. Using a synthetic data set that assumes independent and identically distributed item sets, its performance ratio against the offline optimum converges to one empirically as the number of item sets increases. This work started at a summer internship at HP Labs and resulted in building an automated sponsored search bidding agent that is currently used to place bids for an online advertiser. An outstanding research question in this area is to derive theoretical results about the optimality of the algorithm. As another direction for future work I want to investigate the dynamics of the sponsored search game when all agents follow the proposed strategy.

The proliferation of the Internet brought into existence new ways of sharing information (e.g., BitTorrent), social interaction (e.g., Facebook), shopping (e.g., eBay, Orbitz), and advertising (e.g., AdWords), just to name a few. These systems are used by self-interested participants (e.g., shoppers, advertisers). Consequently, designing a good system or a good strategy for the participants requires game-theoretic considerations. In my future research, I plan to continue to apply economic theory and techniques from computer science to design optimal mechanisms and compute equilibrium strategies for existing and emerging domains.