

Mechanism Design for Stochastic Optimization Problems (Research Overview)

SAMUEL IEONG and MUKUND SUNDARARAJAN

Stanford University

and

ANTHONY MAN-CHO SO

Chinese University of Hong Kong

We identify and address algorithmic and game-theoretic issues arising from welfare maximization in the well-studied two-stage stochastic optimization framework. In contrast, prior work in algorithmic mechanism design has focused almost exclusively on optimization problems without uncertainty. We show both positive results, by demonstrating a mechanism that implements the social welfare maximizer in *sequential ex post equilibrium*, and also negative results, by showing the impossibility of dominant-strategy implementation. In this letter, we describe the relationship between mechanism design and stochastic optimization, and highlight our key technical results. An extended abstract will appear in WINE 2007, and a journal version is under preparation.

Categories and Subject Descriptors: J.4 [Social and Behavioral Sciences]: Economics; F.2.0 [Analysis of Algorithms and Problem Complexity]: General

General Terms: Algorithms, Mechanism Design

Additional Key Words and Phrases: Stochastic Programming

1. INTRODUCTION

Decision makers often have to choose among alternatives in the presence of uncertainty. One popular framework that models such situations is the two-stage stochastic optimization framework [Dantzig 1955; Immorlica et al. 2004; Shmoys and Swamy 2006; Swamy and Shmoys 2006]. In this framework, the decision-making process is separated into two stages. In the first stage, given a probability distribution over possible problem instances (called *scenarios*), the decision maker deploys some resources and incurs some cost. Typically, such an initial deployment is not a feasible solution to every possible scenario, but represents a *hedge* on her part. In the second stage, once a specific scenario is realized, she takes *recourse* actions that augment her initial solution to ensure feasibility, and incurs an additional cost for doing so. The goal of the decision maker is to minimize her *expected* cost (or maximize her expected profit).

Authors' email addresses: {sieurong,mukunds,mancho}@cs.stanford.edu. Samuel Ieong is supported by a Stanford Graduate Fellowship and NSF grant ITR-0205633.

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 2007 ACM 1529-3785/2007/0700-0001 \$5.00

Researchers in stochastic optimization focus on developing efficient algorithms for finding (approximately) good solutions to the problem. These algorithms, however, generally assume that the distribution of the scenarios is known (or that there exists a sampling oracle). In practice, the distribution is often only known to private parties, who have their own objectives. In order for the decision maker to, say, find the socially optimal solution, she needs to elicit these private preferences. Such elicitation problems are studied in the field of mechanism design.

Much work to-date in mechanism design has focused on deterministic, single-round settings. Our work focuses on mechanisms that elicit preferences in a more general setting motivated by the stochastic optimization framework. The primary difference is that communication now takes place over two stages. This presents a novel challenge, as mechanisms must now consider the possibility of elaborate lies over multiple stages.

Compared to some recent works on mechanism design in similar multi-stage settings [Bergemann and Välimäki 2006; Cavallo et al. 2006], our work considers the additional challenge of *algorithmically* implementing our proposed mechanisms. As most stochastic optimization problems do not admit polynomial-time solutions, we investigate the impact of the use of approximation algorithms to our incentive guarantees. Due to space constraint, we will not be able to cover our results in detail here, but details can be found in the extended abstract.

2. MAIN TECHNICAL RESULTS

We start by describing our two-stage mechanism design setting. A mechanism consists of four rules—a decision rule and a payment rule for each of the first and the second stages. In the first stage, before a scenario is realized, agents report distributions over their types to the mechanism. The mechanism applies the first-stage decision rule and first-stage payment rule to compute an initial outcome and payments for each agent. A scenario is then realized; the agents now report a realization of their types to the mechanism. It then chooses a recourse action that augments the first-stage outcome by applying the second-stage decision rule and computes second-stage payments using the second-stage payment rule.

A mechanism implements the social welfare maximizing outcome in a given solution concept if (1) the first-stage decision rule picks an outcome that maximizes the expected social welfare given the distribution of scenarios, and the second-stage decision rule picks a recourse action that maximizes the social welfare for the realized scenario and (2) truth-telling by all agents constitutes an equilibrium under that solution concept. The solution concept that we use is *sequential ex-post equilibrium*; under this concept, even if agents know the distributions (but not the realizations) of the other agents' types, reporting their true type distributions is an equilibrium; in the second stage, even if agents know of the other agents' realized types, reporting their true type is an equilibrium. Our main result is:

THEOREM 2.1. *There exists a family of mechanisms that implement the social welfare maximizer in sequential ex post equilibrium.*¹

¹The second stage decision rule and payment rules resemble VCG; in fact we can show that truth-telling is a dominant strategy in the second stage. The first stage payment rule can be computed

We now briefly discuss why sequential ex-post equilibria are appropriate for our mechanism design setting. In a two-stage setting, when agents' space of possible types is sufficiently rich, for any mechanism that maximizes welfare, there exist scenarios where agents "regret" their first-stage reports. This motivates the use of a *sequential* equilibrium concept, that, in each stage, quantifies only over information available *up to that stage*. For instance, contrast an *ex post equilibrium* with its sequential counterpart defined above. The former would require that truthful first stage reports to be an equilibrium even when agents know the future realization of types of other agents.

Further, we can show that for any welfare maximizing mechanism that satisfies some additional technical conditions (no positive transfers and voluntary participation) truth-telling cannot be a dominant strategy in the first stage. Specifically, when some agents report according to a different distribution in the second stage than the one reported in the first stage (such an agent must not be truth-telling in at least one of the two stages), we can show that it is no longer a best-response for other agents to truth-tell in the first stage.

3. FOR MORE ...

Readers interested in more details are referred to an extended abstract on this research that will appear in WINE 2007. In the paper, we also introduce a novel generalization of the Fixed-Tree Multicast problem [Feigenbaum et al. 2001], and develop an additive approximation algorithm for it. We then investigate to what degree one can preserve incentive guarantees in a mechanism using this algorithm, and characterize a precise trade-off between the running time of the algorithm and the incentive guarantees of the mechanism.

ACKNOWLEDGMENTS

We would like to thank David Parkes and Bob Wilson for fruitful discussion, and Bob in particular for suggesting the term *sequential ex post equilibrium*.

REFERENCES

- BERGEMANN, D. AND VÄLIMÄKI, J. 2006. Efficient dynamic auctions. Working paper.
- CAVALLO, R., PARKES, D. C., AND SINGH, S. 2006. Optimal coordinated planning amongst self-interested agents with private state. In *Proceedings of UAI 2006*.
- DANTZIG, G. B. 1955. Linear programming under uncertainty. *Mgmt. Sci.* 1, 3/4, 197–206.
- FEIGENBAUM, J., PAPADIMITRIOU, C. H., AND SHENKER, S. 2001. Sharing the cost of multicast transmissions. *Journal of Computer and System Sciences* 63, 1, 21–41.
- IMMORLICA, N., KARGER, D., MINKOFF, M., AND MIRROKNI, V. S. 2004. On the costs and benefits of procrastination: approximation algorithms for stochastic combinatorial optimization problems. In *SODA '04*. 691–700.
- SHMOYS, D. B. AND SWAMY, C. 2006. An approximation scheme for stochastic linear programming and its application to stochastic integer programs. *JACM* 53, 6, 978–1012.
- SWAMY, C. AND SHMOYS, D. B. 2006. Approximation algorithms for 2-stage stochastic optimization problems. *SIGACT News* 37, 1, 33–46.

by using backward induction to carefully align incentives.