The Complexity of Simplicity in Mechanism Design

AVIAD RUBINSTEIN
UC Berkeley

Optimal mechanisms are often prohibitively complicated, leading to serious obstacles both in theory and in bridging theory and practice. Consider the problem of a monopolist seller facing a single additive buyer with independent valuations over n heterogeneous items. Even in this simple setting, it is known that optimal (revenue-maximizing) mechanisms may require randomization [Hart and Reny 2012], use menus of infinite size [Daskalakis et al. 2015], and may be computationally intractable [Daskalakis et al. 2014].

In a letter here last year, Babiaoff et al. [Babaioff et al. 2014a] described their attempt to alleviate the problem by showing that a constant fraction of the optimal revenue can be obtained by a simple mechanism. In this letter we argue in favor of a related research direction: finding the *optimal simple mechanism.* We survey our recent results in this setting [Rubinstein 2016] and draw attention to the question of *what is a "simple" mechanism?*

Categories and Subject Descriptors: J.4 [**Social and Behavioral Science**]: Economics

General Terms: Algorithms; Economics, Theory

Additional Key Words and Phrases: Optimal Mechanisms, Simple Mechanisms, Revenue, Approximation

1. INTRODUCTION

Mechanism Design 101. Perhaps the most central problem in mechanism design is this: find the revenue-maximizing mechanism for selling n items to a single buyer with an additive valuation function. Without any constraints, the optimum mechanism is obvious: The buyer is forced to give you all her money. Of course we don't like this solution concept (and neither does the buyer...), so we add the constraint of individual rationality (IR): the buyer will never pay more than her utility from the items allocated. The optimum IR mechanism is also trivial: Charge the buyer her valuation. The problem here is that the true valuation is the buyer's private information and you need incentives to make her reveal it – this is the incentive compatibility (IC) constraint.

The third constraint. So we want to find the optimal IR and IC mechanism. That's much harder to do, but recent works on this problem provide insights into the optimal mechanisms for some important special cases. We now know that those mechanisms require randomization [Hart and Reny 2012], must use a menu of infinite size [Daskalakis et al. 2015], and may be computationally intractable [Daskalakis et al. 2014]. Using these mechanisms is problematic for a variety of reasons: Buyers and sellers may be reluctant to participate in mechanisms that are too complicated (let alone computationally intractable or infinite). Randomization may be restricted by law, or difficult to implement in a trustworthy way, and further complicated by our poor understanding of risk aversion. Such mechanisms

Authors' addresses: aviad@eecs.berkeley.edu

are just as unrealistic – and arguably as unfair to the buyer – as robbing her money or mind-reading her valuation. We need a **simplicity constraint** (SC): the mechanism should be simple. So, the question becomes:

Can we find the optimal IR, IC, and simple mechanism?

Alas, it is not clear how to formalize "simple". In fact, we argue that there is no satisfactory universal definition. "Simple" *can and should* mean different things in different settings; for example, compare the simplicity desiderata in the following scenarios: selling produce in a grocery store (buyers are limited in time and computational capacity); spectrum auctions (buyers may be limited by legal constraints); and ad-auctions in an online marketplace (decisions are often made by automated algorithms). Nevertheless, we believe that it is an important open question to identify useful definitions of "simple", even for special cases of interest.

Alternatively, one may hope to avoid this inconvenient question by showing that "obviously simple" mechanisms approximate the optimal revenue well. In particular Babiaoff et al.'s recent letter [Babaioff et al. 2014a], describes a line of works [Hart and Nisan 2012; Li and Yao 2013; Babaioff et al. 2014b] in our setting that culminated with a celebrated 1/6-approximation of the optimal revenue by the better of the following two mechanisms: (a) sell each item separately; and (b) auction all the items together as one grand bundle. As a corollary, this mechanism gives a 1/6-approximation to the optimal simple mechanism for *any* reasonable definition of "simple".

Our results. We argue that for the important special case of a monopolist facing a single buyer with additive, independent valuations, *partition mechanisms* are a good candidate for a standard for simplicity. (In a partition mechanisms the seller partitions the set of items into disjoint bundles, and posts a price for each bundle; the buyer is allowed to select any number of bundles.)

For this class of mechanisms, our technical contributions include a PTAS, i.e. for any constant $\delta > 0$, we give a polynomial time algorithm that finds a partition mechanism that obtains $(1 - \delta)$-approximation to the optimal revenue among all partition mechanisms. Rather than developing novel algorithmic techniques, our main tool is exploring the structural properties of near-optimal partitions. For example, we prove that there exists a near-optimal partition mechanism with only a constant number of non-trivial bundles. We also prove that this problem is strongly NP-hard, i.e. there is no FPTAS (assuming $\mathsf{P} \neq \mathsf{NP}$).

2. PARTITION MECHANISMS AS SIMPLE MECHANISMS

In this section we briefly argue for the merits of partition mechanisms as a standard for simple mechanisms in our particular setting. There are also some important disadvantages - see further discussion in [Rubinstein 2016].

Expressiveness. We argue that despite their simplicity, partition mechanisms can be used to express important auctions of interest. For example, they generalize both selling items separately and bundling all the items together; thus by [Babaioff et al. 2014b] they guarantee at least a 1/6-approximation to the optimal revenue achievable with any mechanism. Furthermore, this is a strong generalization: as we show in [Rubinstein 2016], partition mechanisms can obtain twice the revenue obtained by the better of selling items separately or bundling all the items together.

Menu complexity and false-name-proofness. Hart and Nisan [Hart and Nisan 2013] discuss a measure of menu-size complexity: every truthful mechanism can be represented as a menu of (potentially randomized) outcomes and prices, where the buyer is allowed to choose one of those outcomes. As noted by Hart and Nisan, the mechanism which auctions each item separately has exponential menu-size complexity under this definition. To overcome this problem, they also introduce a measure that they call *additive-menu-size*, where the buyer is allowed to buy an arbitrary number of outcomes from the menu. Under this definition, partition auctions have linear additive-menu-size complexity.

A related issue is that of false-name-proofness, i.e. can a buyer gain from participating in the mechanism several times? Partition mechanisms (and additive-menu mechanisms in general) have the advantage that they are always false-name-proof.

Locality and buyer-side computational complexity. Partition mechanisms also have the advantage that the buyer's decisions are "local", i.e. the decision to buy one bundle is independent of the decision to buy other bundles. This greatly simplifies tasks such as analyzing and reasoning about such mechanisms, learning or predicting the effects of changes to the environment or the mechanism, etc. In particular, this makes the buyer's decisions very easy.

REFERENCES

Babaioff, M., Immorlica, N., Lucier, B., and Weinberg, S. M. 2014a. A simple and approximately optimal mechanism for an additive buyer. *SIGecom Exchanges 13,* 2, 31–35.

Babaioff, M., Immorlica, N., Lucier, B., and Weinberg, S. M. 2014b. A simple and approximately optimal mechanism for an additive buyer. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*. 21–30.

Daskalakis, C., Deckelbaum, A., and Tzamos, C. 2014. The Complexity of Optimal Mechanism Design. In *the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.

Daskalakis, C., Deckelbaum, A., and Tzamos, C. 2015. Strong duality for a multiple-good monopolist. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015*. 449–450.

Hart, S. and Nisan, N. 2012. Approximate revenue maximization with multiple items. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. EC '12. 656–656.

Hart, S. and Nisan, N. 2013. The menu-size complexity of auctions. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*. EC '13. 565–566.

Hart, S. and Reny, P. J. 2012. Maximal Revenue with Multiple Goods: Nonmonotonicity and Other Observations. Discussion Paper Series dp630, The Center for the Study of Rationality, Hebrew University, Jerusalem. Nov.

Li, X. and Yao, A. C.-C. 2013. On revenue maximization for selling multiple independently distributed items. *Proceedings of the National Academy of Sciences 110,* 28, 11232–11237.

Rubinstein, A. 2016. On the computational complexity of optimal simple mechanisms. In *ITCS*. To appear. Available online at http://arxiv.org/abs/1511.04741.