Are Two (Samples) Really Better Than One?

On the Non-Asymptotic Performance of Empirical Revenue Maximization

MOSHE BABAI OFF, Microsoft Research

YANNAI A. GONCZAROWSKI, The Hebrew University of Jerusalem and Microsoft Research

YISHAY MANSOUR, Tel Aviv University and Google Research

SHAY MORAN, Institute for Advanced Study

The literature on “mechanism design from samples,” which has flourished in recent years at the interface of economics and computer science, offers a bridge between the classic computer-science approach of worst-case analysis (corresponding to “no samples”) and the classic economic approach of average-case analysis for a given Bayesian prior (conceptually corresponding to the number of samples tending to infinity). Nonetheless, the two directions studied so far are two extreme and almost diametrically opposed directions: that of asymptotic results where the number of samples grows large, and that where only a single sample is available. In this paper, we take a first step toward understanding the middle ground that bridges these two approaches: that of a fixed number of samples greater than one. In a variety of contexts, we ask what is possibly the most fundamental question in this direction: *are two samples really better than one sample?* We present a few surprising negative results, and complement them with our main result: showing that the worst-case, over all regular distributions, expected-revenue guarantee of the *Empirical Revenue Maximization* algorithm given two samples is greater than that of this algorithm given one sample. The proof is technically challenging, and provides the first result that shows that *some* deterministic mechanism constructed using two samples can guarantee more than one half of the optimal revenue.

CCS Concepts: •Theory of computation → Design and analysis of algorithms; Sample complexity and generalization bounds; Algorithmic game theory; Algorithmic mechanism design; Convergence and learning in games; Computational pricing and auctions;

Additional Key Words and Phrases: Revenue maximization; empirical revenue maximization; samples

A full version of this paper is available at <https://arxiv.org/abs/1802.08037>.

Authors’ email addresses: Babai off: moshe@microsoft.com; Gonczarowski: yannai@gonch.name; Mansour: mansour@tau.ac.il; Moran: shaymoran1@gmail.com.

The research was done while Mansour and Moran were co-affiliated with Microsoft Research.

Yannai Gonczarowski is supported by the Adams Fellowship Program of the Israel Academy of Sciences and Humanities; his work is supported by ISF grant 1435/14 administered by the Israeli Academy of Sciences, by Israel-USA Bi-national Science Foundation (BSF) grant number 2014389, and by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 740282). The research of Yishay Mansour was supported in part by a grant from the Israel Science Foundation, a grant from United States-Israel Binational Science Foundation (BSF), and the Israeli Centers of Research Excellence (I-CORE) program (Center No. 4/11). The research of Shay Moran is supported by the National Science Foundation under agreement No. CCF-1412958 and by the Simons Foundations. We thank Zhiyi Huang and Tim Roughgarden for stimulating conversations.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

ACM EC'18, June 18–22, 2018, Ithaca, NY, USA.

© 2018 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-4529-3/18/06.

<http://dx.doi.org/10.1145/3219166.3219187>