

A General Theory of Sample Complexity for Multi-Item Profit Maximization

MARIA-FLORINA BALCAN, Carnegie Mellon University, USA

TUOMAS SANDHOLM, Carnegie Mellon University, USA

ELLEN VITERCIK, Carnegie Mellon University, USA

One of the most tantalizing and long-standing open problems in mechanism design is profit maximization in multi-item, multi-buyer settings. Much of the literature surrounding this problem rests on the strong assumption that the mechanism designer knows the distribution over buyers' values. In reality, this information is rarely available. The support of the distribution alone is often doubly exponential, so obtaining and storing the distribution is impractical. We relax this assumption and instead assume that the mechanism designer only has a set of samples from the distribution. We develop learning-theoretic foundations of sample-based mechanism design. In particular, we provide *generalization guarantees* which bound the difference between the empirical profit of a mechanism over a set of samples and its expected profit on the unknown distribution.

In this paper, we present a general theory for deriving worst-case generalization bounds in multi-item settings, as well as data-dependent guarantees when the distribution over buyers' values is well-behaved. We analyze mechanism classes that have not yet been studied in the sample-based mechanism design literature and match or improve over the best-known guarantees for many of the special classes that have been studied. The classes we study include randomized mechanisms, pricing mechanisms, multi-part tariffs, and generalized VCG auctions such as affine maximizer auctions. In the following, we enumerate our main contributions.

A general theory of worst-case generalization guarantees for profit maximization. We uncover a key structural property shared by many mechanisms which allows us to prove strong generalization guarantees: for any fixed set of bids, profit is a piecewise linear function of the mechanism's parameters. Our main theorem provides generalization guarantees for any class exhibiting this structure. To prove this theorem, we relate the complexity of the partition splitting the parameter space into linear portions to the intrinsic complexity of the mechanism class, which we quantify via *pseudo-dimension*. In turn, pseudo-dimension bounds imply generalization bounds. We prove that many mechanisms throughout economics, artificial intelligence, and theoretical computer science share this structure, and thus our main theorem yields strong guarantees.

We prove that our main theorem applies to randomized mechanisms, making us the first to provide generalization bounds for these mechanisms. Our guarantees apply to lotteries, a general representation of randomized mechanisms. Our results imply, for example, that if the mechanism designer plans to offer a menu of ℓ lotteries over m items to an additive or unit-demand buyer, then $\tilde{O}(U^2 \ell m / \epsilon^2)$ samples are sufficient to ensure that every menu's expected profit is ϵ -close to its empirical profit, where U is the maximum profit achievable over the support of the buyer's valuation distribution.

We also provide guarantees for pricing mechanisms using our main theorem. These include item-pricing mechanisms, also known as *posted-price mechanisms*, where each item has a price and buyers buy their

This work was supported in part by the National Science Foundation under grants CCF-1422910, CCF-1535967, IIS-1618714, IIS-1718457, IIS-1617590, CCF-1733556, a Microsoft Research Faculty Fellowship, an Amazon Research Award, a NSF Graduate Research Fellowship, a Microsoft Research Women's Fellowship, and the ARO under award W911NF-17-1-0082.

Authors' addresses: Maria-Florina Balcan, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA, 15213, USA, ninamf@cs.cmu.edu; Tuomas Sandholm, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA, 15213, USA, sandholm@cs.cmu.edu; Ellen Vitercik, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA, 15213, USA, vitercik@cs.cmu.edu.

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).

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

ACM EC'18, June 18–22, 2018, Ithaca, NY, USA. ACM ISBN 978-1-4503-4529-3/18/06.

<https://doi.org/10.1145/3219166.3219217>

utility-maximizing bundles. Additionally, we study *multi-part tariffs*, where there is an upfront fee and a price per unit. We are the first to provide generalization bounds for these tariffs and other non-linear pricing mechanisms, which have been studied in economics for decades. For instance, our main theorem guarantees that if there are κ units of a single good for sale, then $\tilde{O}(U^2\kappa/\epsilon^2)$ samples are sufficient to learn a nearly optimal two-part tariff.

Our main theorem implies generalization bounds for many auction classes, including second price auctions and generalized VCG auctions, such as affine maximizer auctions, virtual valuations combinatorial auctions, and mixed-bundling auctions, which have been studied in artificial intelligence and economics.

A key challenge which differentiates our generalization guarantees from those typically found in machine learning is the sensitivity of these mechanisms to small changes in their parameters. For example, changing the price of a good can cause a steep drop in profit if the buyer no longer wants to buy it. Meanwhile, for many well-understood function classes in machine learning, there is a close connection between the distance in parameter space between two parameter vectors and the distance in function space between the two corresponding functions. Understanding this connection is often the key to quantifying the class's intrinsic complexity. Intrinsic complexity typically translates to pseudo-dimension or another metric which allows us to derive learnability guarantees. Since profit functions do not exhibit this predictable behavior, we must carefully analyze the structure of the mechanisms we study in order to derive strong generalization guarantees.

Data-dependent generalization guarantees for profit maximization. We provide several data-dependent tools that strengthen our main theorem when the distribution over buyers' values is "well-behaved." We prove generalization guarantees that, surprisingly, are independent of the number of items for item-pricing mechanisms, second price auctions with reserves, and a subset of lottery mechanisms. Under anonymous prices, our bounds do not depend on the number of bidders either. These guarantees hold when the bidders are additive with values drawn from *item-independent distributions* (bidder i_1 's value for item j is independent from her value for item j' , but her value for item j may be correlated with bidder i_2 's value for item j).

We also provide data-dependent generalization guarantees that, at a high level, are robust to outliers. Worst-case generalization guarantees typically grow linearly with the maximum profit achievable over the support of the distribution. They are thus pessimistic when the highest valuations have a low probability mass. We show how to obtain stronger guarantees in this setting.

Structural profit maximization. Many of the mechanism classes we study exhibit a hierarchical structure. For example, when designing a pricing mechanism, the designer can segment the population into k groups and charge each group a different price. This is prevalent throughout daily life: movie theaters and amusement parks have different admission prices per market segment, with groups such as Child, Student, Adult, and Senior Citizen. In the simplest case, $k = 1$ and the prices are anonymous. If k equals the number of buyers, the prices are non-anonymous. In general, the designer should not choose the simplest class to optimize over simply to guarantee good generalization because more complex classes are more likely to contain nearly optimal mechanisms. We show how the mechanism designer can determine the precise level in the hierarchy assuring him the optimal tradeoff between profit maximization and generalization.

In summary, by uncovering structure shared by a myriad of pricing, auction, and lottery mechanisms, we prove strong sample complexity bounds for mechanism classes not yet studied in the sample-based mechanism design literature and match or improve over the best known guarantees for many classes. We strengthen our main results with data-dependent bounds when the distribution over buyers' values is "well-behaved." Finally, we investigate a fundamental tradeoff in sample-based mechanism design: complex mechanisms often have higher profit than simple mechanisms, but more samples are required to ensure that empirical and expected profit are close. We provide techniques for optimizing this tradeoff.

Full Paper: <https://arxiv.org/pdf/1705.00243.pdf>.

CCS Concepts: • Theory of computation → Sample complexity and generalization bounds; Computational pricing and auctions;

Additional Key Words and Phrases: Learning theory, machine learning, profit maximization, revenue maximization, combinatorial auctions.