

Mechanism Design, Machine Learning, and Pricing Problems

MARIA-FLORINA BALCAN

and

AVRIM BLUM

School of Computer Science, Carnegie Mellon University

Consider a seller with multiple digital goods or services for sale, such as movies, software, or network services, over which buyers may have complicated preferences. In order to sell these items through an incentive-compatible auction mechanism, this mechanism should have the property that each bidder is offered a set of prices that do not depend on the value of her bid. The problem of designing a revenue-maximizing auction is known in the economics literature as the optimal auction design problem [Myerson 1981]. The classical model for optimal auction design assumes a Bayesian setting in which players' valuations (types) are drawn from some probability distribution that furthermore is known to the mechanism designer. For example, to sell a single item of fixed marginal cost, one should set the price that maximizes the profit margin per sale times the probability a random person would be willing to buy at that price. However, in complex or non-static environments, these assumptions become unrealistic. In these settings, machine learning can provide a natural approach to the design of near-optimal mechanisms without such strong assumptions or degree of prior knowledge.

Specifically, notice that while a truthful auction mechanism should have the property that the prices offered to some bidder i do not depend on the value of her bid, they can depend on the amounts bid by other bidders j . From a Machine Learning perspective, this is very similar to thinking of bidders as "examples" and our objective being to use information from examples $j \neq i$ to produce a good prediction with respect to example i . Thus, without presuming a known distribution over bidders (or even that bidders come from any distribution at all) perhaps if the number of bidders is sufficiently large, enough information can be learned from some of them to perform well on the rest. In recent work [Balcan et al. 2007] we formalize this idea and show indeed that sample-complexity techniques from machine learning theory [Anthony and Bartlett 1999; Vapnik 1998] can be adapted to this setting to give quantitative bounds for this kind of approach. More generally, we show that sample complexity analysis can be applied to convert incentive-compatible mechanism design problems to more standard algorithm-design questions, in a wide variety of revenue-maximizing auction settings.

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

Our reductions imply that for these problems, given an algorithm for the non incentive-compatible pricing problem, we can convert it into an algorithm for the incentive-compatible mechanism design problem that is only a factor of $(1 + \epsilon)$ worse, so long as the number of bidders is sufficiently large as a function of an appropriate measure of complexity of the class of allowable pricings. We apply these results to the problem of auctioning a digital good [Goldberg et al. 2001; Goldberg et al. 2006], to the attribute auction problem which includes a wide variety of discriminatory pricing problems [Blum and Hartline 2005; Aggarwal and Hartline 2006], and to the problem of item-pricing in unlimited-supply combinatorial auctions [Guruswami et al. 2005]. From a machine learning perspective, these settings present several challenges: in particular, the *loss function* is discontinuous, is asymmetric, and has a large range.

The high level idea of the most basic reduction in [Balcan et al. 2007] which is based on the idea of a random sampling auction is actually fairly natural. For concreteness, let us imagine we are selling a collection of n goods or services of zero marginal cost to us, to m bidders who may have complex preference functions over these items, and our objective is to achieve revenue comparable to the best possible assignment of prices to the various items we are selling. Thus, we are in the setting of maximizing revenue in an unlimited supply combinatorial auction. Then given a set of bids S , we perform the following operations. We first randomly partition S into two sets S_1 and S_2 . We then consider the purely algorithmic problem of finding the best set of prices p_1 for the set of bids S_1 (which may be difficult but is purely algorithmic), and the best set of prices p_2 for the set of bids S_2 . We then use p_1 as offer prices for bidders in S_2 , giving each bidder the bundle maximizing revealed valuation minus price, and use p_2 as offer prices for bidders in S_1 . We then show that even if bidders' preferences are extremely complicated, this mechanism will achieve revenue close to that of the best fixed assignment of prices to items so long as the number of bidders is sufficiently large compared to the number of items for sale. For example, if all bidders' valuations on the grand bundle of all n items lie in the range $[1, h]$, then $O(hm/\epsilon^2)$ bidders are sufficient so that with high probability, we come within a $(1 + \epsilon)$ factor of the optimal fixed item pricing. Or, if we cannot solve the algorithmic problem exactly (since many problems of this form are often NP-hard [Guruswami et al. 2005; Balcan and Blum 2007; Balcan et al. 2007; Briest and Krysta 2006]), we lose only a $(1 + \epsilon)$ factor over whatever approximation our method for solving the algorithmic problem gives us.

More generally, these methods apply to a wide variety of pricing problems, including those in which bidders have both public and private information, and also give a formal framework in which one can address other interesting design issues such as how fine-grained a market segmentation should be. This framework provides a unified approach to considering a variety of profit maximizing mechanism design problems including many that have been previously considered in the literature. Furthermore, our results substantially generalize the previous work on random sampling mechanisms by both broadening the applicability of such mechanisms and by simplifying the analysis.

Some of our techniques give suggestions for the design of mechanisms and others for their analysis. In terms of design, these include the use of discretization to produce smaller function classes, and the use of structural-risk minimization to choose an appropriate level of complexity of the mechanism for a given set of bidders. In terms of analysis, these include both the use of basic sample-complexity arguments, and the notion of multiplicative covers for better bounding the true complexity of a given class of offers.

Finally, from a learning perspective, this mechanism-design setting presents a number of technical challenges when attempting to get good bounds: in particular, the payoff function is discontinuous and asymmetric, and the payoffs for different offers are non-uniform. For example, we develop bounds based on a different notion of covering number than typically used in machine learning, in order to obtain results that are more meaningful for this mechanism design setting.

Acknowledgements. This work was supported in part by the NSF grant CCF-0514922.

REFERENCES

- AGGARWAL, G. AND HARTLINE, J. 2006. Knapsack Auctions. In *Proceedings of the 17th ACM-SIAM Symposium on Discrete Algorithms*. 1083 – 1092.
- ANTHONY, M. AND BARTLETT, P. 1999. *Neural Network Learning: Theoretical Foundations*. Cambridge University Press.
- BALCAN, M.-F. AND BLUM, A. 2007. Approximation Algorithms and Online Mechanisms for Item Pricing. *Theory of Computing* 3, 179 – 195. Originally appeared in the The Seventh ACM Conference on Electronic Commerce (EC), 2006.
- BALCAN, M.-F., BLUM, A., CHAN, H., AND HAJIAGHAYI, M. 2007. A theory of loss-leaders: Making money by pricing below cost. In *Proc. 3rd International Workshop on Internet and Network Economics*. Lecture Notes in Computer Science.
- BALCAN, M.-F., BLUM, A., HARTLINE, J., AND MANSOUR, Y. 2007. Reducing Mechanism Design to Algorithm Design via Machine Learning. *Journal of Computer and System Sciences, special issue on Learning Theory*. Originally appeared in the 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS) 2005.
- BLUM, A. AND HARTLINE, J. 2005. Near-Optimal Online Auctions. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms*. 1156 – 1163.
- BRIEST, P. AND KRZYSTA, P. 2006. Single-minded unlimited supply pricing on sparse instances. In *Proc. 17th Annual ACM-SIAM Symposium on Discrete Algorithms*. 1093–1102.
- GOLDBERG, A., HARTLINE, J., KARLIN, A., SAKS, M., AND WRIGHT, A. 2006. Competitive Auctions and Digital Goods. *Games and Economic Behavior*.
- GOLDBERG, A., HARTLINE, J., AND WRIGHT, A. 2001. Competitive auctions and digital goods. In *Proc. 12th Annual ACM-SIAM Symposium on Discrete Algorithms*. 735–744.
- GURUSWAMI, V., HARTLINE, J., KARLIN, A., KEMPE, D., KENYON, C., AND MCSHERRY, F. 2005. On Profit-Maximizing Envy-Free Pricing. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms*. 1164 – 1173.
- MYERSON, R. 1981. Optimal Auction Design. *Mathematics of Operations Research* 6, 58–73.
- VAPNIK, V. 1998. *Statistical Learning Theory*. Springer-Verlag.