
Bayesian Mechanism Design

Bayesian Mechanism Design

Jason D. Hartline

Northwestern University

Evanston, 60208

USA

hartline@eecs.northwestern.edu

now

the essence of **knowledge**

Boston – Delft

Foundations and Trends[®] in Theoretical Computer Science

Published, sold and distributed by:

now Publishers Inc.
PO Box 1024
Hanover, MA 02339
USA
Tel. +1-781-985-4510
www.nowpublishers.com
sales@nowpublishers.com

Outside North America:

now Publishers Inc.
PO Box 179
2600 AD Delft
The Netherlands
Tel. +31-6-51115274

The preferred citation for this publication is J. D. Hartline, Bayesian Mechanism Design, Foundation and Trends[®] in Theoretical Computer Science, vol 8, no 3, pp 143–263, 2012

ISBN: 978-1-60198-670-2

© 2013 J. D. Hartline

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, mechanical, photocopying, recording or otherwise, without prior written permission of the publishers.

Photocopying. In the USA: This journal is registered at the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923. Authorization to photocopy items for internal or personal use, or the internal or personal use of specific clients, is granted by now Publishers Inc. for users registered with the Copyright Clearance Center (CCC). The 'services' for users can be found on the internet at: www.copyright.com

For those organizations that have been granted a photocopy license, a separate system of payment has been arranged. Authorization does not extend to other kinds of copying, such as that for general distribution, for advertising or promotional purposes, for creating new collective works, or for resale. In the rest of the world: Permission to photocopy must be obtained from the copyright owner. Please apply to now Publishers Inc., PO Box 1024, Hanover, MA 02339, USA; Tel. +1-781-871-0245; www.nowpublishers.com; sales@nowpublishers.com

now Publishers Inc. has an exclusive license to publish this material worldwide. Permission to use this content must be obtained from the copyright license holder. Please apply to now Publishers, PO Box 179, 2600 AD Delft, The Netherlands, www.nowpublishers.com; e-mail: sales@nowpublishers.com

**Foundations and Trends[®] in
Theoretical Computer Science**
Volume 8 Issue 3, 2012
Editorial Board

Editor-in-Chief:

Madhu Sudan

*Microsoft Research New England
One Memorial Drive
Cambridge, Massachusetts 02142
USA*

Editors

Bernard Chazelle (Princeton)
Oded Goldreich (Weizmann Inst.)
Shafi Goldwasser (MIT and Weizmann Inst.)
Sanjeev Khanna (University of Pennsylvania)
Jon Kleinberg (Cornell University)
László Lovász (Microsoft Research)
Christos Papadimitriou (UC. Berkeley)
Prabhakar Raghavan (Yahoo! Research)
Peter Shor (MIT)
Madhu Sudan (Microsoft Research)
Éva Tardos (Cornell University)
Avi Wigderson (IAS)

Editorial Scope

Foundations and Trends[®] in Theoretical Computer Science

will publish survey and tutorial articles in the following topics:

- Algorithmic game theory
- Computational algebra
- Computational aspects of combinatorics and graph theory
- Computational aspects of communication
- Computational biology
- Computational complexity
- Computational geometry
- Computational learning
- Computational Models and Complexity
- Computational Number Theory
- Cryptography and information security
- Data structures
- Database theory
- Design and analysis of algorithms
- Distributed computing
- Information retrieval
- Operations Research
- Parallel algorithms
- Quantum Computation
- Randomness in Computation

Information for Librarians

Foundations and Trends[®] in Theoretical Computer Science, 2012, Volume 8, 4 issues. ISSN paper version 1551-305X. ISSN online version 1551-3068. Also available as a combined paper and online subscription.

Foundations and Trends[®] in
Theoretical Computer Science
Vol. 8, No. 3 (2012) 143–263
© 2013 J. D. Hartline
DOI: 10.1561/04000000045



Bayesian Mechanism Design

Jason D. Hartline

*Northwestern University, Evanston, 60208, USA,
hartline@eecs.northwestern.edu*

Abstract

Systems wherein strategic agents compete for limited resources are ubiquitous: the economy, computer networks, social networks, congestion networks, nature, etc. Assuming the agents' preferences are drawn from a distribution, which is a reasonable assumption for small mechanisms in a large system, *Bayesian mechanism design* governs the design and analysis of these systems.

This article surveys the classical economic theory of Bayesian mechanism design and recent advances from the perspective of algorithms and approximation. Classical economics gives simple characterizations of Bayes-Nash equilibrium and optimal mechanisms when the agents' preferences are linear and single-dimensional. The mechanisms it predicts are often complex and overly dependent on details of the model. Approximation complements this theory and suggests that simple and less-detail-dependent mechanisms can be nearly optimal. Furthermore, techniques from approximation and algorithms can be used to describe good mechanisms beyond the single-dimensional, linear model of agent preferences.

Contents

1	Introduction	1
1.1	Topics Covered	4
1.2	Topics Omitted	8
2	Equilibrium	11
2.1	Equilibrium	13
2.2	Independent, Single-dimensional, and Linear Utilities	16
2.3	Equilibrium Characterization	17
2.4	Solving for Equilibrium	19
3	Optimal Mechanisms	21
3.1	Single-dimensional Environments	22
3.2	Single-agent Revenue Maximization	25
3.3	Multi-agent Revenue Maximization	33
3.4	Regular Distributions and Ironing	35
4	Approximation Mechanisms	41
4.1	Reserve Pricing	42
4.2	Posted Pricing	45
4.3	Prior-independent Approximation	54

5 Multi-dimensional and Non-linear Preferences	61
5.1 Single-agent Optimization	62
5.2 Multi-agent Optimization with Revenue Linearity	67
5.3 Multi-agent Optimization without Revenue Linearity	71
5.4 Multi-agent Optimization with Multi-dimensional Externalities	80
6 Approximation for Multi-dimensional and Non-linear Preferences	85
6.1 Single-agent Approximation	85
6.2 Multi-agent Approximation in Service Constrained Environments	90
6.3 Multi-agent Approximation with Multi-dimensional Externalities	96
7 Computation and Approximation Algorithms	101
7.1 Single-dimensional Monotonization via Ironing	103
7.2 Multi-dimensional Monotonization via Matching	107
A Mathematical Reference	111
A.1 Submodular Set Functions	111
A.2 Matroid Set Systems	111
A.3 Convex Optimization	113
Acknowledgments	117
References	119

1

Introduction

The Bayesian¹ approach to optimization assumes that an input to the optimization problem is drawn from a distribution and requests that an output be found that is good in expectation. This approach is both compatible with the standard worst-case approach to algorithm design and suggests a rich problem space that is relatively unexplored from an algorithmic perspective. The approach is to partition the optimization problem into two stages. In the first stage, a distribution is given and can be preprocessed so as to construct an algorithm tailored to the distribution. In the second stage inputs to the distribution are drawn and the algorithm is run on these inputs. For instance, this approach is similar to data structure problems such as Huffman [64] coding and search problems such as nearest neighbor, cf., Cover and Hart [36].

¹We say “Bayesian” here instead of “stochastic” to connote settings where *Bayesian updating* is relevant. As an example, if the distribution over inputs is correlated and an optimization algorithm must make decisions before all the input is known, the algorithm should perform Bayesian updating to refine its beliefs about remaining inputs. That said, many of the methods and results described in this survey pertain to the special case that the inputs are independent and the term “Bayesian” is only used for consistency in presenting the result in the context of the greater literature within which they are contained.

2 Introduction

Furthermore, the algorithm practitioner's design and analysis framework can be abstracted as being Bayesian. As a caricature, such a practitioner will design and test her algorithm on a class of data sets that she deems relevant. When the attempted improvements to the algorithm on these data sets fail to significantly increase performance, she will deploy the final algorithm. Methodologies for the Bayesian framework for algorithm design can then potentially be applied to this form of practical algorithm design (see, e.g., the *Bayesian reductions* of Section 7).

This survey focuses on a large class of Bayesian optimization problems that exhibit both algorithmic and economic challenges. In particular, in a departure from traditional algorithms, we assume that the desired input to the algorithm is the private information of self-interested agents. These agents can report this information to the algorithm (if requested) but can equally well misreport if it benefits them and the algorithm cannot tell the difference. Our abstract question pertains to how a designer should structure the rules of a system so that in the equilibrium of strategic play, the designer's objective is optimized. This is the question of *mechanism design*.

Mechanism design has broad applications; however, we will focus on potential applications to computer systems. In so far as scarce system resources must be shared amongst parties with diverse and selfish interests, mechanism design governs the proper working of computer systems. When incentives are not properly accounted for, strategic misbehavior is common. A few examples include spam in email systems (see Dwork et al. [44]), link-spam and Sybil networks in the ranking of Internet search results (see Dwork et al. [45] and Gyongyi and Garcia-Molina [55]), and freeloading in file-sharing networks (see Vishnumurthy et al. [92]).

The literature that melds algorithmic and economic issues in mechanism design, a.k.a., *algorithmic mechanism design*, since its inception over a decade ago, has predominantly focused on worst-case mechanisms with dominant strategy equilibria. The strategic agents should have a single best action no matter what the actions of other agents, and the designed mechanism should be good for any preferences the agents may have. As a motivating example, the second-price auction

for selling a single item serves the highest bidder at price equal to the second highest bid. Truth-telling is a dominant strategy in the second-price auction; therefore, in dominant strategy equilibrium the bidder with the highest true value wins and, for any values the bidders might possess, the auction maximizes social welfare. In this line of worst-case dominant-strategy mechanism design, Lehmann et al. [70] initiated the merger of algorithmic complexity and economic incentive considerations for the objective of social welfare maximization, Nisan and Ronen [75] initiated the study of mechanisms for non-linear objectives such as makespan in machine scheduling, and Goldberg et al. [50] initiated the study of objectives that depend on agent payments (e.g., revenue). The latter two agendas are related in that, even absent computational constraints, the economic incentives of the agents preclude a pointwise optimal mechanism. As an example, notice that for the single-item auction, a seller would could obtain more revenue by setting a reserve price between the highest and second highest value, but to do so requires knowledge of the agent values. Therefore, for these problems the worst-case optimization question under consideration is inherently one of approximation akin to the competitive analysis of online algorithms (e.g., Borodin and El-Yaniv [16]).

In the classic microeconomic treatment of mechanism design, the non-pointwise-optimality of mechanisms is resolved by formulating the problem as one of Bayesian design. A distribution over the preferences of the agents is given, and the designer seeks to optimize her objective in expectation over this distribution. For any specific distribution, there is such an optimal mechanism. The Bayesian assumption also allows the strategic incentives of the agents to be relaxed. Given that the agents' preferences are drawn from a distribution, instead of requiring a mechanism to have dominant strategies, mechanisms can be considered where agents' actions are best responses to the distribution of actions of other agents.

Both the Bayesian and worse-case (henceforth, *prior-free*) frameworks for mechanism design have merits and, recognizing this, a Bayesian branch of algorithmic mechanism design has emerged. Paramount of study in *Bayesian algorithmic mechanism design* are algorithmic techniques, approximation, and computational issues. The

4 Introduction

goal of this survey is to discuss the most fundamental of these results in the context of classical Bayesian mechanism design.

1.1 Topics Covered

This survey is organized into the following sections. For environments with agents with linear and single-dimensional preferences, Section 2 characterizes the equilibrium of strategic play in auction-like games, Section 3 characterizes optimal auctions for welfare and revenue, and Section 4 describes simple, practical mechanism that are approximately optimal. For environments where agents have multi-dimensional and non-linear preferences, Section 5 characterizes optimal auctions and Section 6 describes simple approximation mechanisms. Finally, Section 7 gives generic performance-preserving reductions from mechanism design to algorithm design for both single-dimensional and multi-dimensional agent preferences. Mathematical reference is given in Appendix A.

Section 2: Equilibrium. In Bayesian games an agent's strategy maps her private information to an action in the game; the distribution of private information and strategies induce a distribution of actions in the game; and a profile of strategies is in equilibrium if the strategy of each agent is a best response to this distribution of actions. The classical Bayesian approach to mechanism design starts with a characterization of equilibrium. For single-dimensional environments, i.e., when an agent's private preference is given by a single number denoting her value for a single abstract service, Myerson [74] characterized all possible equilibria. This characterization states that an agent's probability of service should be monotonically non-decreasing in her value and it gives a formula for her expected payment. This payment identity implies the *revenue equivalence* of auctions with the same equilibrium outcome. Moreover, it suggests a method for solving for equilibrium.

Section 3: Optimal Mechanisms. For welfare maximization, the well known Vickrey-Clarke-Groves (VCG) mechanism is pointwise optimal [91, 35, 53]. For revenue maximization in single-dimensional

environments, Myerson [74] gives a reduction from the non-pointwise problem of maximizing revenue in expectation over the distribution, to the problem of optimizing a pointwise virtual welfare. The reduction works by transforming the agent values to *virtual values*, and then optimizing in the transformed space. In the special case where values are i.i.d. from a distribution from a sufficiently well-behaved class, this optimal mechanism is simply the second-price auction with a suitably chosen reserve price. Bulow and Roberts [20] give a microeconomic reinterpretation Myerson's virtual values as the derivative of an appropriate *revenue curve*, a.k.a., as a *marginal revenue*. Alaei et al. [5] observe that this characterization is based on the "revenue linearity" of optimal mechanisms for single-dimensional agents.

Section 4: Approximation Mechanisms. While the optimal mechanism is often complex and impractical, there is often a simple and practical mechanism that is approximately optimal. For example, the *second-price auction with reserve* is widely prevalent even though it is not optimal beyond the ideal setting of symmetric agents with values drawn from a well-behaved distribution. Of course, even more prevalent are simple posted-pricing mechanisms, e.g., most stores sell goods by posting take-it-or-leave-it while-supplies-last prices on the goods for sale. Approximation can resolve this disconnect between theory and practice. Chawla et al. [29, 30] and Hartline and Roughgarden [62] show that reserve pricing is approximately optimal in many environments. Similarly, Chawla et al. [30], Yan [94], and Chakraborty et al. [27] show that posted pricings (take-it-or-leave-it while-supplies-last prices) are approximately optimal quite broadly.

As discussed previously, for many mechanism design problems there is not a single optimal mechanism. For a given prior distribution, the optimal mechanism generally depends on the prior distribution. Nonetheless, there may still be a single *prior-independent* mechanism that is approximately optimal. I.e., for any distribution, the prior-independent mechanism approximates the optimal mechanism for that distribution. Dhangwatnotai et al. [41] show that a single sample from the distribution gives a sufficient market analysis for obtaining a two approximation to the revenue-optimal auction in

6 Introduction

many single-dimensional environments; moreover, a single sample can be easily attained on-the-fly as the mechanism is being run. Hartline and Roughgarden [61] provide a post hoc Bayesian justification of the preceding literature on prior-free revenue approximation (e.g., Goldberg et al. [49]) and show that prior-free approximation with respect to an appropriate prior-free benchmark implies prior-independent approximation.

Section 5: Multi-dimensional and Non-linear Preferences.

Revenue maximization in multi-dimensional environments is much more complex than in single-dimensional environments. Nonetheless, a series of recent papers has given an algorithmic generalization of the single-dimensional reduction to virtual welfare maximization of Myerson [74] to multi-dimensional preferences. Cai et al. [22] and Alaei et al. [4] do so for environments where welfare is optimized by a greedy algorithm; and Cai et al. [23] extend these results to general environments with additive preferences. These results tie the theory of optimal mechanism design to a natural convex optimization problem. While these approaches result in mechanisms that have polynomial complexity in the number of agents, they rely on brute-force solutions to single-agent pricing problems.

An important special case of multi-dimensional mechanism design is single-agent pricing; the unit-demand pricing problem is a paradigmatic challenge problem. Consider a single agent who desires one of a set of items. This agent is multi-dimensional in that she may have a distinct value for each item. The agent's multi-dimensional preference is drawn from a distribution and, given this distribution, we would like to price items to maximize revenue. Briest et al. [18] show that this problem can be solved in time polynomial in the number of distinct agent types. Unfortunately, when the agent's values for the items are independent, the type space is exponentially big in the number of items.

Section 6: Approximation for Multi-dimensional and Non-linear Preferences.

The aforementioned unit-demand pricing problem with independently distributed values can be simplified with approximation. Chawla et al. [29, 30] show that there is a simple two

approximation to the optimal item pricing. Moreover, Chawla et al. [32] show that this two approximation to the optimal item pricing is in fact also a four approximation to the optimal mechanism which, in addition to pricing items, may also price lotteries over items. Chawla et al. [30, 32], Bhattacharya et al. [11], Chakraborty et al. [27], and Alaei [2] extend these results quite broadly to show that often for unit-demand auction problems simple to find posted-pricing-based mechanisms are approximately optimal.

As described above, in both single-dimensional and multi-dimensional revenue maximization, the (non-pointwise) Bayesian mechanism design problem reduces to a (pointwise) virtual welfare maximization problem. It should be noted, however, that the single-dimensional case and multi-dimensional case have rather different structure. In the single-dimensional case, the transformation from value space to virtual value space is deterministic and separates across agents; whereas, in the multi-dimensional case, the transformation requires solving a convex optimization problem on all the agents together and it may be stochastic. Alaei et al. [5] show that, in fact, there is a transformation for the multi-dimensional case, with similar structure and economic intuition as in the single-dimensional case, that is approximately optimal.

Section 7: Computation and Approximation Algorithms. It is standard (from the prior-free mechanism design literature) that there is a reduction from exact welfare maximization with incentives to exact welfare maximization without incentives. I.e., if we have an optimal algorithm for maximizing welfare, we can convert that algorithm into an optimal mechanism that, in the dominant-strategy equilibrium of strategic play, maximizes welfare. The resulting mechanism is known as the Vickrey-Clarke-Groves (VCG) mechanism [35, 53, 91]. Lehmann et al. [70] point out, however, this reduction is incompatible with generic approximation algorithms. I.e., from a generic approximation algorithm we cannot instantiate the reduction to obtain a generic approximation mechanism which has an equilibrium with performance comparable to the original algorithm.

8 Introduction

While mounting evidence suggests that such generic approximation-compatible reductions (see Chawla et al. [31] and Dobzinski and Vondrák [42]) do not exist for prior-free mechanisms, they do for Bayesian mechanisms. For welfare maximization in single-dimensional environments, Hartline and Lucier [60] give a generic approximation-preserving reduction from Bayesian mechanism design to Bayesian algorithm design. (Notice that, of course, a worst-case algorithm is also a Bayesian algorithm.) Via the reduction from revenue to welfare maximization of Myerson [74], this reduction can be adapted to the revenue objective. Hartline et al. [59] and Bei and Huang [9] generalize the above single-dimensional reduction to multi-dimensional environments. The multi-dimensional approach is brute-force in each agent's type space and it is an open question as to whether a similar reduction exists for large but succinctly represented type spaces.

Appendix A: Mathematical Reference. A number of mathematical constructs play a prominent role in our treatment of Bayesian mechanism design. These are *submodular set functions* which capture the concept of diminishing returns, *matroid set systems* which represent substitutability (e.g., Oxley [76]), and *convex optimization* within which most questions in Bayesian mechanism design reside (e.g., Schrijver [86]).

1.2 Topics Omitted

Having described above the material covered by this survey, we now turn to related material that is not covered. We have omitted discussion of almost all of the literature on prior-free mechanism design. Prior-free mechanism design and recent results relating to computational tractability and approximation are a topic warranting a survey of their own, which to cover adequately, would be even longer than this survey.

Also notably absent from this survey is discussion of non-revelation mechanisms, i.e., mechanisms that do not have truth-telling or otherwise easy-to-find equilibria. The big challenge of analyzing non-revelation mechanisms is that almost any departure from simple symmetric environments renders solving for equilibrium analytically intractable

(cf. Section 2.4). To address this challenge, methodologies from the literature on quantifying the *price of anarchy*, i.e., the suboptimality of performance that results from strategic behavior, can be employed. Instead of explicitly solving for equilibrium, a price-of-anarchy analysis considers minimal necessary properties of equilibria to argue that any equilibrium must be pretty good. Of course, this would be a problematic exercise if we believed that, just as we cannot solve for equilibrium, neither can the agents. Fortunately, because the analysis employs only minimal assumptions, a frequent corollary of a price-of-anarchy analysis is that best-response dynamics and no-regret learning algorithms have good performance even if they never reach an equilibrium.

To illustrate the power of the price-of-anarchy approach, we will describe a few recent results. Consider the welfare objective in multi-agent multi-item environment (which we consider in Section 5 and Section 6 for the revenue objective). Suppose instead of running a simultaneous auction that coordinates the sale of the items, we run independent first- or second-price auctions either simultaneously or sequentially. These auctions are strategically complex, e.g., because there is an exposure problem where an agent may win multiple items even if she only wants one. A series of papers, Bikhchandani [13], Christodoulou et al. [34], Bhawalkar and Roughgarden [12], Hassidim et al. [63], Paes Leme and Tardos [78], Paes Leme et al. [79], Roughgarden [83], and Syrgkanis and Tardos [88, 89] showed that in various configurations the price of anarchy of independent auctions for multiple items is often a constant like two. Paes Leme et al. [77] give a short survey of these results.

Price of anarchy approaches have also been applied to an auction known as the *generalized second-price* auction which is used by Internet search engines to sell advertisements that are displayed alongside search results (see, e.g., Fain and Pedersen [46]). In this auction, bidders are ranked by bid and each bidder is charged the bid of her successor. This auction does not have simple equilibria as does the second-price auction. Gomes and Sweeney [51] show that even in symmetric Bayesian settings the generalized-second-price auction may not have any efficient equilibria; subsequent work of Caragiannis et al. [25] bounded the potential inefficiency of its equilibria by a factor of slightly under three.

10 *Introduction*

A final class of results looks at non-revelation mechanisms based on approximation algorithms (because, as mentioned above, there are no general approaches for converting approximation algorithms into mechanisms without Bayesian assumptions). Lucier and Borodin [71] consider multi-dimensional combinatorial auctions based on greedy approximation algorithms. They show that in equilibrium, a mechanism based on a greedy β approximation algorithm is at worst a $(\beta + 1)$ approximation in equilibrium.

The results described above are primarily for the welfare objective, and methodologies from the price of anarchy have seen relatively less success for the revenue objective. One exception is by Lucier et al. [72] who give revenue bounds for the previously discussed generalized-second-price auction.

References

- [1] S. Agrawal, Y. Ding, A. Saberi, and Y. Ye, “Price of correlations in stochastic optimization,” *Operations Research*, vol. 60, no. 1, pp. 150–162, 2012.
- [2] S. Alaei, “Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers,” in *Proceedings of the 52nd Annual Symposium on Foundations of Computer Science*, pp. 512–521, IEEE, 2011.
- [3] S. Alaei, “Mechanism design with general utilities,” PhD thesis, University of Maryland, 2012.
- [4] S. Alaei, H. Fu, N. Haghpanah, J. Hartline, and A. Malekian, “Bayesian optimal auctions via multi-to single-agent reduction,” in *Proceedings of the ACM Conference on Electronic Commerce*, pp. 17–17, 2012. Available at: <http://arxiv.org/abs/1203.5099>.
- [5] S. Alaei, H. Fu, N. Haghpanah, J. Hartline, and A. Malekian, “The simple economics of approximately optimal auctions,” arXiv preprint arXiv:1206.3541, 2012.
- [6] M. Armstrong, “Price discrimination by a many-product firm,” *Review of Economic Studies*, vol. 66, no. 1, pp. 151–168, 1999.
- [7] M.-F. Balcan, A. Blum, J. Hartline, and Y. Mansour, “Reducing mechanism design to algorithm design via machine learning,” *Journal of Computer and System Sciences*, vol. 74, no. 8, pp. 1245–1270, 2008.
- [8] S. Baliga and R. Vohra, “Market research and market design,” *Advances in Theoretical Economics*, vol. 3, 2003.
- [9] X. Bei and Z. Huang, “Bayesian incentive compatibility via fractional assignments,” in *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 720–733, SIAM, 2011.

120 *References*

- [10] D. Bergemann and S. Morris, “Robust mechanism design,” *Econometrica*, vol. 73, pp. 1771–1813, 2005.
- [11] S. Bhattacharya, G. Goel, S. Gollapudi, and K. Munagala, “Budget constrained auctions with heterogeneous items,” in *Proceedings of the 42nd ACM symposium on Theory of computing*, pp. 379–388, ACM, 2010.
- [12] K. Bhawalkar and T. Roughgarden, “Welfare guarantees for combinatorial auctions with item bidding,” in *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 700–709, SIAM, 2011.
- [13] S. Bikhchandani, “Auctions of heterogeneous objects,” *Games and Economic Behavior*, vol. 26, no. 2, pp. 193–220, 1999.
- [14] K. Border, “Implementation of reduced form auctions: A geometric approach,” *Econometrica*, vol. 59, no. 4, pp. 1175–1187, 1991.
- [15] K. Border, “Reduced form auctions revisited,” *Economic Theory*, vol. 31, no. 1, pp. 167–181, 2007.
- [16] A. Borodin and R. El-Yaniv, *Online Computation and Competitive Analysis*. Cambridge University Press, 1998.
- [17] P. Briest, “Uniform budgets and the envy-free pricing problem,” in *Automata, Languages and Programming*, pp. 808–819, Springer, 2008.
- [18] P. Briest, S. Chawla, R. Kleinberg, and M. Weinberg, “Pricing randomized allocations,” in *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 585–597, Society for Industrial and Applied Mathematics, 2010.
- [19] J. Bulow and P. Klemperer, “Auctions versus negotiations,” *American Economic Review*, vol. 86, pp. 180–194, 1996.
- [20] J. Bulow and J. Roberts, “The simple economics of optimal auctions,” *The Journal of Political Economy*, vol. 97, pp. 1060–1090, 1989.
- [21] Y. Cai and C. Daskalakis, “Extreme-value theorems for optimal multidimensional pricing,” in *Proceedings of the 52nd Annual Symposium on Foundations of Computer Science*, pp. 522–531, IEEE, 2011.
- [22] Y. Cai, C. Daskalakis, and M. Weinberg, “An algorithmic characterization of multi-dimensional mechanisms,” in *Symposium on Theory of Computation*, pp. 459–478, 2012.
- [23] Y. Cai, C. Daskalakis, and M. Weinberg, “Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization,” in *Foundations of Computer Science*, pp. 130–139, 2012.
- [24] B. Caillaud and J. Robert, “Implementation of the revenue-maximizing auction by an ignorant seller,” *Review of Economic Design*, vol. 9, no. 2, pp. 127–143, 2005.
- [25] I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, M. Kyropoulou, B. Lucier, R. Paes Leme, and É. Tardos, “On the efficiency of equilibria in generalized second price auctions,” arXiv preprint arXiv:1201.6429, 2012.
- [26] C. Carathéodory, “Über den Variabilitätsbereich der Fourierschen Konstanten von positiven harmonischen Funktionen,” *Rendiconti del Circolo Matematico di Palermo (1884–1940)*, vol. 32, no. 1, pp. 193–217, 1911.
- [27] T. Chakraborty, E. Even-Dar, S. Guha, Y. Mansour, and S. Muthukrishnan, “Approximation schemes for sequential posted pricing in multi-unit auctions,” *Internet and Network Economics*, pp. 158–169, 2010.

- [28] S. Chawla and J. Hartline, “Auctions with unique equilibria,” in *Proceedings of the ACM Conference on Electronic Commerce*, 2013.
- [29] S. Chawla, J. Hartline, and R. Kleinberg, “Algorithmic pricing via virtual valuations,” in *Proceedings of the ACM Conference on Electronic Commerce*, 2007.
- [30] S. Chawla, J. Hartline, D. Malec, and B. Sivan, “Sequential posted pricing and multi-parameter mechanism design,” in *Proceedings of the 41st ACM Symposium on Theory of Computing*, 2010.
- [31] S. Chawla, N. Immorlica, and B. Lucier, “On the limits of black-box reductions in mechanism design,” in *Proceedings of the 44th symposium on Theory of Computing*, pp. 435–448, ACM, 2012.
- [32] S. Chawla, D. Malec, and B. Sivan, “The power of randomness in bayesian optimal mechanism design,” in *Proceedings of the ACM Conference on Electronic Commerce*, pp. 149–158, 2010.
- [33] B. Chazelle, “An optimal convex hull algorithm in any fixed dimension,” *Discrete & Computational Geometry*, vol. 10, no. 1, pp. 377–409, 1993.
- [34] G. Christodoulou, A. Kovács, and M. Schapira, “Bayesian combinatorial auctions,” *Automata, Languages and Programming*, pp. 820–832, 2008.
- [35] E. H. Clarke, “Multipart Pricing of Public Goods,” *Public Choice*, vol. 11, pp. 17–33, 1971.
- [36] T. Cover and P. Hart, “Nearest neighbor pattern classification,” *Information Theory, IEEE Transactions on*, vol. 13, no. 1, pp. 21–27, 1967.
- [37] N. Devanur, B. Ha, and J. Hartline, “Prior-free auctions for budgeted agents,” in *Proceedings of the ACM Conference on Electronic Commerce*, 2013.
- [38] N. Devanur and J. Hartline, “Limited and online supply and the bayesian foundations of prior-free mechanism design,” in *Proceedings of the ACM Conference on Electronic Commerce*, pp. 41–50, 2009.
- [39] N. Devanur, J. Hartline, A. Karlin, and T. Nguyen, “Prior-independent multi-parameter mechanism design,” *Internet and Network Economics*, pp. 122–133, 2011.
- [40] N. Devanur, J. Hartline, and Q. Yan, “Envy freedom and prior-free mechanism design,” *CoRR*, abs/1212.3741, 2012.
- [41] P. Dhangwatnotai, T. Roughgarden, and Q. Yan, “Revenue maximization with a single sample,” in *ECOM10*, 2010.
- [42] S. Dobzinski and J. Vondrák, “The computational complexity of truthfulness in combinatorial auctions,” in *Proceedings of the ACM Conference on Electronic Commerce*, pp. 405–422, ACM, 2012.
- [43] S. Dughmi and T. Roughgarden, “Black-box randomized reductions in algorithmic mechanism design,” in *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science*, pp. 775–784, IEEE, 2010.
- [44] C. Dwork, A. Goldberg, and M. Naor, “On memory-bound functions for fighting spam,” *Advances in Cryptology-Crypto 2003*, pp. 426–444, 2003.
- [45] C. Dwork, R. Kumar, M. Naor, and D. Sivakumar, “Rank aggregation methods for the web,” in *Proceedings of the 10th international conference on World Wide Web*, pp. 613–622, ACM, 2001.
- [46] D. Fain and J. Pedersen, “Sponsored search: A brief history,” *Bulletin of the American Society for Information Science and Technology*, vol. 32, no. 2, pp. 12–13, 2006.

122 *References*

- [47] H. Fu, J. Hartline, and D. Hoy, “Prior-independent Auctions for Risk-averse Agents,” in *Proceedings of the ACM Conference on Electronic Commerce*, 2013.
- [48] A. Goldberg and J. Hartline, “Collusion-Resistant Mechanisms for Single-Parameter Agents,” in *Proceedings of the 16th ACM Symposium on Discrete Algorithms*, 2005.
- [49] A. Goldberg, J. Hartline, A. Karlin, M. Saks, and A. Wright, “Competitive Auctions,” *Games and Economic Behavior*, vol. 55, pp. 242–269, 2006.
- [50] A. Goldberg, J. Hartline, and A. Wright, “Competitive auctions and digital goods,” in *Proceedings of the 12th ACM Symposium on Discrete Algorithms*, pp. 735–744, ACM/SIAM, 2001.
- [51] R. Gomes and K. Sweeney, “Bayes-Nash equilibria of the generalized second-price auction,” *Games and Economic Behavior*, 2012.
- [52] M. Grötschel, L. Lovász, and A. Schrijver, “The ellipsoid method and its consequences in combinatorial optimization,” *Combinatorica*, vol. 1, no. 2, pp. 169–197, 1981.
- [53] T. Groves, “Incentives in Teams,” *Econometrica*, vol. 41, pp. 617–631, 1973.
- [54] F. Gul and E. Stacchetti, “Walrasian equilibrium with gross substitutes,” *Journal of Economic Theory*, vol. 87, no. 1, pp. 95–124, 1999.
- [55] Z. Gyongyi and H. Garcia-Molina, “Web spam taxonomy,” in *First international workshop on adversarial information retrieval on the web (AIRWeb 2005)*, 2005.
- [56] N. Haghpanah, “Optimal multi-dimensional pricings,” manuscript, 2013.
- [57] G. Hardy, J. Littlewood, and G. Pólya, “Some simple inequalities satisfied by convex functions,” *Messenger of Math*, vol. 58, pp. 145–152, 1929.
- [58] J. Hartline and A. Karlin, “Profit maximization in mechanism design,” in *Algorithmic Game Theory*, ch. 13, (N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, eds.), pp. 331–362, Cambridge University Press, 2007.
- [59] J. Hartline, R. Kleinberg, and A. Malekian, “Bayesian incentive compatibility via matchings,” in *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 734–747, SIAM, 2011.
- [60] J. Hartline and B. Lucier, “Bayesian Algorithmic Mechanism Design,” in *Proceedings of the 41st ACM Symposium on Theory of Computing*, 2010.
- [61] J. Hartline and T. Roughgarden, “Optimal Mechanism Design and Money Burning,” in *Proceedings of the 39th ACM Symposium on Theory of Computing*, 2008.
- [62] J. Hartline and T. Roughgarden, “Simple versus Optimal Mechanisms,” in *Proceedings of the 10th ACM Conference on Electronic Commerce*, 2009.
- [63] A. Hassidim, H. Kaplan, Y. Mansour, and N. Nisan, “Non-price equilibria in markets of discrete goods,” in *Proceedings of the ACM Conference on Electronic Commerce*, pp. 295–296, New York, NY, USA: ACM, 2011.
- [64] D. Huffman, “A method for the construction of minimum-redundancy codes,” *Proceedings of the IRE*, vol. 40, no. 9, pp. 1098–1101, 1952.
- [65] M. Jackson and H. Sonnenschein, “Overcoming incentive constraints by linking decisions,” *Econometrica*, vol. 75, no. 1, pp. 241–257, 2007.
- [66] R. Kleinberg and M. Weinberg, “Matroid prophet inequalities,” in *Proceedings of the 44th symposium on Theory of Computing*, pp. 123–136, ACM, 2012.

- [67] J. Kruskal, “On the shortest spanning subtree of a graph and the traveling salesman problem,” in *Proceedings of the American Mathematical Society*, vol. 7, pp. 48–50, 1956.
- [68] J.-J. Laffont and J. Robert, “Optimal auction with financially constrained buyers,” *Economics Letters*, vol. 52, no. 2, pp. 181–186, 1996.
- [69] R. Lavi and C. Swamy, “Truthful and Near-Optimal Mechanism Design via Linear Programming,” in *Proceedings of the 46th IEEE Symposium on Foundations of Computer Science*, 2005.
- [70] D. Lehmann, L. O’Callaghan, and Y. Shoham, “Truth revelation in approximately efficient combinatorial auctions,” *Journal of the ACM (JACM)*, vol. 49, no. 5, pp. 577–602, 2002.
- [71] B. Lucier and A. Borodin, “Price of anarchy for greedy auctions,” in *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, Society for Industrial and Applied Mathematics, pp. 537–553, 2010.
- [72] B. Lucier, R. Paes Leme, and É. Tardos, “On revenue in the generalized second price auction,” in *Proceedings of the 21st international conference on World Wide Web*, pp. 361–370, ACM, 2012.
- [73] S. Matthews, “Selling to risk averse buyers with unobservable tastes,” *Journal of Economic Theory*, vol. 30, no. 2, pp. 370–400, August 1983.
- [74] R. Myerson, “Optimal Auction Design,” *Mathematics of Operations Research*, vol. 6, pp. 58–73, 1981.
- [75] N. Nisan and A. Ronen, “Algorithmic Mechanism Design,” *Games and Economic Behavior*, vol. 35, pp. 166–196, 2001.
- [76] J. Oxley, *Matroid Theory*. Oxford Science Publications. The Clarendon Press, Oxford University Press, New York, 1992.
- [77] R. Paes Leme, V. Syrgkanis, and E. Tardos, “The Dining Bidder Problem: a la russe et la française,” *SIGECOM Exchanges*, pp. 25–28, 2012.
- [78] R. Paes Leme, V. Syrgkanis, and É. Tardos, “Sequential auctions and externalities,” in *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 869–886, SIAM, 2012.
- [79] R. Paes Leme and E. Tardos, “Pure and Bayes-Nash price of anarchy for generalized second price auction,” in *Proceedings of the 2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pp. 735–744, IEEE Computer Society, 2010.
- [80] M. Pai and R. Vohra, “Optimal auctions with financially constrained bidders,” Discussion Papers, Northwestern University, August 2008.
- [81] J. Rochet, “The taxation principle and multi-time Hamilton-Jacobi equations,” *Journal of Mathematical Economics*, vol. 14, no. 2, pp. 113–128, 1985.
- [82] J. Rochet, “A necessary and sufficient condition for rationalizability in a quasi-linear context,” *Journal of Mathematical Economics*, vol. 16, no. 2, pp. 191–200, 1987.
- [83] T. Roughgarden, “The price of anarchy in games of incomplete information,” in *Proceedings of the ACM Conference on Electronic Commerce*, pp. 862–879, ACM, 2012.
- [84] T. Roughgarden, I. Talgam-Cohen, and Q. Yan, “Supply-limiting mechanisms,” in *Proceedings the ACM Conference on Electronic Commerce*, pp. 844–861, 2012.

124 *References*

- [85] E. Samuel-Cahn, “Comparison of threshold stop rules and maximum for independent nonnegative random variables,” *The Annals of Probability*, vol. 12, no. 4, pp. 1213–1216, 1984.
- [86] A. Schrijver, *Combinatorial Optimization: Polyhedra and Efficiency (Algorithms and Combinatorics)*. Springer, July 2003.
- [87] I. Segal, “Optimal pricing mechanisms with unknown demand,” *American Economic Review*, 2003.
- [88] V. Syrgkanis and E. Tardos, “Bayesian sequential auctions,” in *Proceedings of the ACM Conference on Electronic Commerce*, pp. 929–944, 2012.
- [89] V. Syrgkanis and E. Tardos, “Composable and efficient mechanisms,” arXiv preprint arXiv:1211.1325, 2012.
- [90] M. Talagrand, “Matching random samples in many dimensions,” *Annals of Applied Probability*, vol. 2, no. 4, pp. 846–856, 1992.
- [91] W. Vickrey, “Counterspeculation, auctions, and competitive sealed tenders,” *Journal of Finance*, vol. 16, pp. 8–37, 1961.
- [92] V. Vishnumurthy, S. Chandrakumar, and E. Sirer, “Karma: A secure economic framework for peer-to-peer resource sharing,” in *Workshop on Economics of Peer-to-Peer Systems*, 2003.
- [93] R. Wilson, *Nonlinear Pricing*. Oxford University Press, 1997.
- [94] Q. Yan, “Mechanism design via correlation gap,” in *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 710–719, SIAM, 2011.