

# **Learning in Repeated Auctions**

**Other titles in Foundations and Trends® in Machine Learning**

*Dynamical Variational Autoencoders: A Comprehensive Review*

Laurent Girin, Simon Leglaive, Xiaoyu Bie, Julien Diard, Thomas Hueber and Xavier Alameda-Pineda

ISBN: 978-1-68083-912-8

*Machine Learning for Automated Theorem Proving: Learning to Solve SAT and QSATe*

Sean B. Holden

ISBN: 978-1-68083-898-5

*Spectral Methods for Data Science: A Statistical Perspective*

Yuxin Chen, Yuejie Chi, Jianqing Fan and Cong Ma

ISBN: 978-1-68083-896-1

*Tensor Regression*

Jiani Liu, Ce Zhu, Zhen Long and Yipeng Liu

ISBN: 978-1-68083-886-2

*Minimum-Distortion Embedding*

Akshay Agrawal, Alnur Ali and Stephen Boyd

ISBN: 978-1-68083-888-6

*Graph Kernels: State-of-the-Art and Future Challenges*

Karsten Borgwardt, Elisabetta Ghisu, Felipe Llinares-López, Leslie O'Bray and Bastian Rieck

ISBN: 978-1-68083-770-4

# Learning in Repeated Auctions

---

**Thomas Nedelec**

ENS Paris Saclay

Criteo AI Lab

thomas.nedelec@polytechnique.edu

**Clément Calauzènes**

Criteo AI Lab

c.calauzenes@criteo.com

**Noureddine El Karoui**

Work done while at

UC, Berkeley

Criteo AI Lab

nkarouiprof@gmail.com

**Vianney Perchet**

ENSAE

Criteo AI Lab

vianney.perchet@normalesup.org

**now**

the essence of knowledge

Boston — Delft

## Foundations and Trends<sup>®</sup> in Machine Learning

*Published, sold and distributed by:*

now Publishers Inc.  
PO Box 1024  
Hanover, MA 02339  
United States  
Tel. +1-781-985-4510  
[www.nowpublishers.com](http://www.nowpublishers.com)  
[sales@nowpublishers.com](mailto: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

T. Nedelec *et al.*. *Learning in Repeated Auctions*. Foundations and Trends<sup>®</sup> in Machine Learning, vol. 15, no. 3, pp. 176–334, 2022.

ISBN: 978-1-68083-939-5

© 2022 T. Nedelec *et al.*

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](http://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](http://www.nowpublishers.com); [sales@nowpublishers.com](mailto: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](http://www.nowpublishers.com); e-mail: [sales@nowpublishers.com](mailto:sales@nowpublishers.com)

# Foundations and Trends<sup>®</sup> in Machine Learning

## Volume 15, Issue 3, 2022

### Editorial Board

#### Editor-in-Chief

**Michael Jordan**

University of California, Berkeley  
United States

#### Editors

Peter Bartlett  
*UC Berkeley*

Yoshua Bengio  
*Université de Montréal*

Avrim Blum  
*Toyota Technological  
Institute*

Craig Boutilier  
*University of Toronto*

Stephen Boyd  
*Stanford University*

Carla Brodley  
*Northeastern University*

Inderjit Dhillon  
*Texas at Austin*

Jerome Friedman  
*Stanford University*

Kenji Fukumizu  
*ISM*

Zoubin Ghahramani  
*Cambridge University*

David Heckerman  
*Amazon*

Tom Heskes  
*Radboud University*

Geoffrey Hinton  
*University of Toronto*

Aapo Hyvarinen  
*Helsinki IIT*

Leslie Pack Kaelbling  
*MIT*

Michael Kearns  
*UPenn*

Daphne Koller  
*Stanford University*

John Lafferty  
*Yale*

Michael Littman  
*Brown University*

Gabor Lugosi  
*Pompeu Fabra*

David Madigan  
*Columbia University*

Pascal Massart  
*Université de Paris-Sud*

Andrew McCallum  
*University of  
Massachusetts Amherst*

Marina Meila  
*University of Washington*

Andrew Moore  
*CMU*

John Platt  
*Microsoft Research*

Luc de Raedt  
*KU Leuven*

Christian Robert  
*Paris-Dauphine*

Sunita Sarawagi  
*IIT Bombay*

Robert Schapire  
*Microsoft Research*

Bernhard Schoelkopf  
*Max Planck Institute*

Richard Sutton  
*University of Alberta*

Larry Wasserman  
*CMU*

Bin Yu  
*UC Berkeley*

## Editorial Scope

### Topics

Foundations and Trends® in Machine Learning publishes survey and tutorial articles in the following topics:

- Adaptive control and signal processing
- Applications and case studies
- Behavioral, cognitive and neural learning
- Bayesian learning
- Classification and prediction
- Clustering
- Data mining
- Dimensionality reduction
- Evaluation
- Game theoretic learning
- Graphical models
- Independent component analysis
- Inductive logic programming
- Kernel methods
- Markov chain Monte Carlo
- Model choice
- Nonparametric methods
- Online learning
- Optimization
- Reinforcement learning
- Relational learning
- Robustness
- Spectral methods
- Statistical learning theory
- Variational inference
- Visualization

### Information for Librarians

Foundations and Trends® in Machine Learning, 2022, Volume 15, 6 issues. ISSN paper version 1935-8237. ISSN online version 1935-8245. Also available as a combined paper and online subscription.

## Contents

---

<b>1</b>	<b>Introduction: Scope and Motivation</b>	<b>3</b>
1.1	Bayesian mechanism design . . . . .	4
1.2	Learning theory and auction design . . . . .	5
1.3	Organization of the monograph . . . . .	8
<b>2</b>	<b>Bayesian Mechanism Design</b>	<b>10</b>
2.1	The Bayesian setting . . . . .	12
2.2	Sealed-bid auctions . . . . .	17
2.3	A revenue equivalence theorem . . . . .	23
2.4	Deriving revenue-maximizing auctions . . . . .	26
2.5	Prior-independent optimal auctions . . . . .	48
2.6	Advanced material: non-unicity of Nash equilibria and related complications . . . . .	52
<b>3</b>	<b>Repeated Auctions From a Seller's Standpoint</b>	<b>58</b>
3.1	Motivations . . . . .	58
3.2	Statistical learning theory tools for revenue maximization . . . . .	59
3.3	Auctions with asymptotically no approximation error . . . . .	69
3.4	Tractability at the cost of approximation error . . . . .	73
3.5	Contextual estimation of reserve prices . . . . .	87
3.6	Cost & online estimation of auctions . . . . .	92

<b>4 Adaptive and Strategic Learning Agents</b>	<b>103</b>
4.1 Adaptive bidders: online learning to bid . . . . .	105
4.2 Mechanism with adaptive bidders & full surplus extraction	115
4.3 Reversing the asymmetry: strategic buyer vs. myopic seller	134
<b>References</b>	<b>149</b>



# Learning in Repeated Auctions

Thomas Nedelec<sup>1</sup>, Clément Calauzènes<sup>2</sup>, Nouredine El Karoui<sup>3</sup> and Vianney Perchet<sup>4</sup>

<sup>1</sup>*ENS Paris Saclay and Criteo AI Lab, France;*  
*thomas.nedelec@polytechnique.edu*

<sup>2</sup>*Criteo AI Lab, France; c.calauzenes@criteo.com*

<sup>3</sup>*Work done while at UC, Berkeley, USA, and Criteo AI Lab, France;*  
*nkarouiprof@gmail.com*

<sup>4</sup>*ENSAE and Criteo AI Lab, France; vianney.perchet@normalesup.org*

---

## ABSTRACT

Online auctions are one of the most fundamental facets of the modern economy and power an industry generating hundreds of billions of dollars a year in revenue. Auction theory has historically focused on the question of designing the best way to sell a single item to potential buyers, with the concurrent objectives of maximizing revenue generated or welfare created. Theoretical results in this area have typically relied on some prior Bayesian knowledge agents were assumed to have on each other. This assumption is no longer satisfied in new markets such as online advertising: similar items are sold repeatedly, and agents are unaware of each other or might try to manipulate each other. On the other hand, statistical learning theory now provides tools to supplement those missing pieces of information given enough data, as agents can learn from their environment to improve their strategies.

This monograph covers recent advances in learning in repeated auctions, starting from the traditional economic

---

Thomas Nedelec, Clément Calauzènes, Nouredine El Karoui and Vianney Perchet (2022), “Learning in Repeated Auctions”, Foundations and Trends<sup>®</sup> in Machine Learning: Vol. 15, No. 3, pp 176–334. DOI: 10.1561/22000000077.

©2022 T. Nedelec *et al.*

study of optimal one-shot auctions with a Bayesian prior. We then focus on the question of learning optimal mechanisms from a dataset of bidders' past values. The sample complexity as well as the computational efficiency of different methods will be studied. We will also investigate online variants where gathering data has a cost to be accounted for, either by sellers or buyers ("earning while learning"). Later in the monograph, we will further assume that bidders are also adaptive to the mechanism as they interact repeatedly with the same seller. We will show how strategic agents can actually manipulate repeated auctions, to their own advantage. A particularly interesting example is that of reserve price improvements for strategic buyers in second price auctions.

All the questions discussed in this monograph are grounded in real-world applications and many of the ideas and algorithms we describe are used every day to power the Internet economy.

---

# 1

---

## Introduction: Scope and Motivation

---

The main purpose of auction theory is to construct a set of rules that will be used by a seller to sell one or several items to a group of potential buyers, who will send messages (or *bids*) to the seller – usually indicating how much they value the item or how much they are willing to pay to acquire it. In almost all cases, it is sufficient to define only two rules. First, the *allocation rule* describes which buyer wins the auction (if a unique non-divisible item is sold), depending on the different messages received; if the item is divisible, the allocation rule describes how the item is shared between winners. Second, the *payment rule* indicates to buyers how much they are going to pay to the seller, again based on the different messages. Those rules are known publicly before the auction starts, and they influence the behavior, or strategy, of the different buyers.

When choosing an allocation and a payment rule, the seller might have several constraints to respect: **1)** maximizing the revenue she is getting from the auction (*revenue maximization*); **2)** ensuring the participation of buyers in the auction and making sure they have an incentive to participate (*individual rationality*); **3)** ensuring that given the rules of the auction, it is in the best interest of buyers to reveal how

much they truly value an item (*incentive compatibility*) as it may make revenue maximization easier. On the other side of the game, the buyers strategically adjust the bids sent to the seller depending on auction rules in order to maximize their own utility.

Historically, auctions have often been designed so that buyers have an incentive to bid in a way that reflects how much they truly value the items that are for sale. This constraint still leaves plenty of choices for auction design, and a large part of the literature has focused on designing auctions that maximize the seller's revenue, assuming buyers are rational. However, with the advent of the Internet and the automation of auctions, the landscape of possible applications has changed drastically, necessitating more complex settings to accurately study the incentives and behaviors at play. More recently, the auction literature has aimed at understanding how the design of an auction platform impacts seller's revenue, and the global welfare and behavior of buyers and sellers in contexts where sellers (and sometimes buyers) participate in a very large number of auctions each day. These setups reflect situations appearing in modern online marketplaces.

### 1.1 Bayesian mechanism design

Auction theory has focused for a long time on the simplest case: there is a single, non-divisible item to be sold to a set of predefined buyers in a one-shot auction. The chosen mechanism indicates which buyer (if any) gets the item and at which price. The seminal works of Vickrey (1961), Myerson (1981) and Riley and Samuelson (1981) emphasize the importance of the information structure of an auction system. It consists of the information owned privately by the buyers and the information that the seller has on each buyer. This information owned privately by the buyers is the value they give to the item, i.e, the highest price they are willing to pay to get the item. The uncertainties upon these different values lie at the gist of the seller's optimization problem: otherwise, she would just have to sell the item to the buyer with the highest value, at this price or infinitesimally less.

To handle this deficit of information about buyers, it is standard to take a "Bayesian" viewpoint and assume that the seller has some

probabilistic prior on the values given to the item by each bidder. This prior distribution is usually called the value distribution and it encompasses the seller's uncertainty on a specific bidder's values. There are of course several possibilities for how this value distribution is constructed. For instance, in wine or art auctions, it often comes from expert knowledge about an admissible price for a good wine bottle or for an important piece of art.

## 1.2 Learning theory and auction design

It is now possible for Internet platforms to run billions of auctions a day and store most of the historical data coming from them. This digitization of auction mechanisms was the first step into gathering data to optimize selling mechanisms. Auctions are now used in most Internet platforms to organize interactions between the different stakeholders. eBay was one of the first big online platforms to use ascending auction to sell objects on the platform. Google and most search engines companies started to use auctions to sell ad opportunities on their front page. For instance, they let advertisers bid on some keywords to get sponsored links above the first results for a certain user query. Nowadays, Facebook and LinkedIn are also using them to determine which ad to display, Amazon and most e-commerce marketplaces decide which products are going to be sponsored (and/or advertised) through an auction mechanism and auctions are also used to sell carbon permits by the European Union or to run large electricity markets.

To exploit this new source of available information (i.e., enormous datasets of past bids), practitioners used advanced statistical learning algorithms in connection with the classical Bayesian theory. Indeed, beyond the AI hype, machine learning algorithms are now widely applied in the industry for numerous applications: the value distribution is no longer coming from some given and fixed prior, but learned (hopefully accurately and efficiently) on historical—bidding—data. The first large-scale field experiment in production showed how engineers at Yahoo could handle their huge datasets to learn an optimal reserve price per key word (Ostrovsky and Schwarz, 2011). This results in data-driven mechanisms whose design use techniques coming from a large

variety of fields, including statistics, machine learning, game theory and Economics. Similarly, bidders on these online platforms also gather data and use new statistical learning techniques to improve their bidding strategies against automated mechanisms. This flood of data and the associated paradigm shift it constitutes opens many new interesting practical problems, new theoretical questions and new interesting games to study.

### 1.2.1 Repeated auctions only from a seller's standpoint

The first natural repeated game setting consists of understanding how the seller can learn a revenue-maximizing auction mechanism from a dataset of bids or values. In the example of eBay marketplace, the seller (eBay) observes numerous auctions a day for similar items. Hence, from its point of view, the mechanism is repeated and she can aim at optimizing some long-term revenue. On the contrary, buyers are individuals that participate in a few, if not a single, auctions at best. Then, from their point of view, the mechanism still looks like a one-shot auction and they are bound to implement myopic short-term strategies, optimizing point-wisely their utility (by opposition to long-term and effectively in expectation). Let us consider the simplifying assumption where bidder values on the platform are sampled from a certain unknown distribution, that encompasses the variability in their readiness to pay a certain price. Assuming the bidders actually bid their true value (for instance, if the mechanism chosen is fixed and “incentive-compatible”, i.e., bidding one’s value is optimal for buyers), the seller has then access at the end of the day to a dataset of buyer values.

Inspired by the computational learning formalism, Elkind (2007), Balcan *et al.* (2008), and Cole and Roughgarden (2014) initiated a line of research aiming at finding approximations of the revenue-maximizing auction, if possible, efficiently, with approximation guarantees depending on the size of the dataset gathered (a.k.a., the sample complexity). This setting is called the *batch learning setting*. A variant considers the case where the flow of buyers is continuously coming on the platform and the seller can update continuously her mechanism. This is *the online*

*learning setting* introduced in Cesa-Bianchi *et al.* (2014). In all these problems, it is crucial that the samples gathered in the dataset have the same distribution as the samples that will be gathered and treated in the future.

### 1.2.2 Repeated auctions from seller and bidder standpoint

The crucial assumption of myopic/short-sighted/impatient bidders facing a patient seller is unfortunately not necessarily satisfied, depending on the setting. In modern-day practice, typically large online ad platforms, such as Google DoubleClick or AppNexus, are selling ad opportunities for large publishers such as some of the biggest online newspapers. The main difference with the aforementioned eBay example is that only a few companies are actually bidding in these auctions. They are furthermore doing so repeatedly and participating in a massive number of auctions.

Indeed, most companies willing to display ads actually rely on third-parties, *demand-side platforms* (DSP), that are buying and displaying ads for them (because of technical constraints, even sending bids in real-time might actually be quite complex). These aggregated bidders are repeatedly interacting with the (same) seller, billions of times a day. Consequently, this type of buyer can also optimize for long-term utility and need not be myopic. Thus, even if the seller is using one-shot incentive compatible auctions - for instance to gather data in order to later design and switch to a revenue maximizing mechanism -, the bidder might have an interest in not bidding “truthfully”, as classical theory would suggest is optimal for them. Indeed, if buyers do not bid their values, this will modify the distribution of “values” observed by the seller. Subsequently, the mechanism chosen to optimize their revenue will be different from what it would have been had the bidders been naïve, to the advantage of the buyers (Tang and Zeng, 2018; Nedelec *et al.*, 2019b).

Intuitively, this is possible because the information asymmetry that arose in the eBay example between the seller and the bidders – one optimizing over the long-term, the other over the short-term – is almost reversed. If the seller must commit to a specific mechanism or a family

of mechanisms, for instance for contractual reasons, and buyers have this information, they can strategically leverage it by, for example, changing their bidding behavior. In the end, the respective utilities of the seller and buyers will somehow depend on the underlying amount of asymmetry between them. Several works have started studying various intermediate settings, for example, when bidders are (almost) identical (Kanoria and Nazerzadeh, 2014), or are patient, but not as patient as the seller (Amin *et al.*, 2013), etc.

### 1.3 Organization of the monograph

In this monograph, our overarching objective is to provide a widely accessible introduction to the fascinating topics of classical and modern auction theory while bringing to the fore the statistical and machine learning lenses to the topic. We will very clearly state the differences between the different information-asymmetry settings we will review, and point to cutting edge theoretical and practical solutions adapted to them. We will also show how new statistical tools can be used to tackle some important and well-known problems from Economics. Furthermore, those questions open many new interesting problems in Economics since algorithms are replacing classical sellers and buyers. We believe that modern auction theory offers a nice framework to understand what data and Computer Science can bring to modern Economics.

In Section 2, we survey the main results of the Bayesian auction literature, initiated with the seminal works of Vickrey and Myerson. Those results form some of the backbone of classical auction theory and are widely used in Internet practice. We will recall what is the revenue-maximizing auction once the seller has a prior on bidder's valuations and introduce some approximations of the revenue-maximizing auction when the seller must use simpler auctions. In Section 3, we focus on the setting derived from the eBay use case and tackle both the batch learning setting and the online learning setting. We recall some key concepts of statistical learning theory, derive the sample complexity of some of the learning algorithms used to compute a revenue-maximizing auction and show their computational complexity. In Section 4, we focus on the less studied but crucially important setting where bidders can



be strategic regarding the mechanism itself since they have multiple interactions with the seller. We review some of the main methods that have been devised to keep bidders from being strategic in that context, show their limitations and introduce some very new results and approaches developed for bidders to take advantage of the seller's learning process.

This monograph only assumes basic familiarity with standard notions of Machine Learning, Statistics and Data Science and is written with a reader having this background in mind. We hope our monograph will be useful to engineers and researchers looking for an introduction to the beautiful and fast developing topics of modern auction theory and applications.

## References

---

- Abeille, M., C. Calauzènes, N. E. Karoui, T. Nedelec, and V. Perchet. (2018). “Explicit shading strategies for repeated truthful auctions”. In: *arXiv preprint arXiv:1805.00256*.
- Agrawal, S., C. Daskalakis, V. S. Mirrokni, and B. Sivan. (2018). “Robust Repeated Auctions under Heterogeneous Buyer Behavior”. In: *Proceedings of the 2018 ACM Conference on Economics and Computation*. 171–171.
- Albert, M., V. Conitzer, and P. Stone. (2017). “Automated design of robust mechanisms”. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 31. No. 1.
- Allouah, A. and O. Besbes. (2020). “Prior-independent optimal auctions”. *Management Science*. 66(10): 4417–4432.
- Amin, K., A. Rostamizadeh, and U. Syed. (2013). “Learning prices for repeated auctions with strategic buyers”. In: *Proceedings of the 26th International Conference on Neural Information Processing Systems-Volume 1*. 1169–1177.
- Amin, K., A. Rostamizadeh, and U. Syed. (2014). “Repeated contextual auctions with strategic buyers”. In: *Proceedings of the 27th International Conference on Neural Information Processing Systems-Volume 1*. 622–630.

- Archer, A. and É. Tardos. (2001). “Truthful mechanisms for one-parameter agents”. In: *Proceedings 2001 IEEE International Conference on Cluster Computing*. IEEE.
- Arlotto, A. and I. Gurvich. (2019). “Uniformly Bounded Regret in the Multisecretary Problem”. *Stochastic Systems*. 9(3): 231–260.
- Armstrong, M. (1996). “Multiproduct nonlinear pricing”. *Econometrica: Journal of the Econometric Society*: 51–75.
- Arnosti, N., M. Beck, and P. Milgrom. (2016). “Adverse selection and auction design for internet display advertising”. *American Economic Review*. 106(10): 2852–66.
- Ashlagi, I., C. Daskalakis, and N. Haghpanah. (2016). “Sequential mechanisms with ex-post participation guarantees”. In: *Proceedings of the 2016 ACM Conference on Economics and Computation*. 213–214.
- Aström, K. J. and R. M. Murray. (2008). *Feedback Systems: An Introduction for Scientists and Engineers*. Princeton University Press.
- Athey, S. and P. A. Haile. (2007). “Chapter 60 Nonparametric Approaches to Auctions”. In: *Handbook of Econometrics*.
- Audibert, J.-Y. and S. Bubeck. (2009). “Minimax policies for adversarial and stochastic bandits”. In: *Proceedings of COLT*.
- Babaioff, M., R. D. Kleinberg, and A. Slivkins. (2010). “Truthful Mechanisms with Implicit Payment Computation”. In: *Proceedings of the 11th ACM Conference on Electronic Commerce. EC '10*. Cambridge, Massachusetts, USA: Association for Computing Machinery. 43–52.
- Babaioff, M., Y. Sharma, and A. Slivkins. (2014). “Characterizing Truthful Multi-armed Bandit Mechanisms”. *SIAM Journal on Computing*. 43(1): 194–230.
- Balcan, M.-F., A. Blum, J. D. Hartline, and Y. Mansour. (2008). “Reducing mechanism design to algorithm design via machine learning”. *Journal of Computer and System Sciences*. 74(8): 1245–1270.
- Balseiro, S. R., O. Besbes, and G. Y. Weintraub. (2015). “Repeated auctions with budgets in ad exchanges: Approximations and design”. *Management Science*. 61(4): 864–884.
- Balseiro, S. R., O. Candogan, and H. Gurkan. (2020). “Multistage Intermediation in Display Advertising”. *Manufacturing & Service Operations Management*.

- Balseiro, S. R. and Y. Gur. (2019). “Learning in repeated auctions with budgets: Regret minimization and equilibrium”. *Management Science*. 65(9): 3952–3968.
- Balseiro, S. R., V. S. Mirrokni, and R. P. Leme. (2018). “Dynamic mechanisms with martingale utilities”. *Management Science*. 64(11): 5062–5082.
- Bar-Yossef, Z., K. Hildrum, and F. Wu. (2002). “Incentive-compatible online auctions for digital goods.” In: *SODA*. Vol. 2. 964–970.
- Bartlett, P. L., S. Boucheron, and G. Lugosi. (2002). “Model selection and error estimation”. *Machine Learning*. 48(1-3): 85–113.
- Blum, A., V. Kumar, A. Rudra, and F. Wu. (2004). “Online learning in online auctions”. *Theoretical Computer Science*. 324(2-3): 137–146.
- Boyd, S. and L. Vandenberghe. (2004). *Convex Optimization*. USA: Cambridge University Press.
- Braverman, M., J. Mao, J. Schneider, and M. Weinberg. (2018). “Selling to a no-regret buyer”. In: *Proceedings of the 2018 ACM Conference on Economics and Computation*. 523–538.
- Bubeck, S. and N. Cesa-Bianchi. (2012). “Regret Analysis of Stochastic and Nonstochastic Multi-armed Bandit Problems”. In: *Machine Learning*. Vol. 5. No. 1. 1–122.
- Bubeck, S., N. R. Devanur, Z. Huang, and R. Niazadeh. (2017). “Multi-scale Online Learning and its Applications to Online Auctions”. *Proceedings of the Eighteenth ACM Conference on Economics and Computation*.
- Bulow, J. and P. Klemperer. (1996). “Auctions Versus Negotiations”. *The American Economic Review*. 86(1): 180–194.
- Celis, L. E., G. Lewis, M. Mobius, and H. Nazerzadeh. (2014). “Buy-it-now or take-a-chance: Price discrimination through randomized auctions”. *Management Science*. 60(12): 2927–2948.
- Cesa-Bianchi, N., T. Cesari, and V. Perchet. (2019). “Dynamic pricing with finitely many unknown valuations”. In: *Algorithmic Learning Theory*. PMLR. 247–273.
- Cesa-Bianchi, N., C. Gentile, and Y. Mansour. (2014). “Regret minimization for reserve prices in second-price auctions”. *IEEE Transactions on Information Theory*. 61(1): 549–564.

- Choi, H., C. F. Mela, S. R. Balseiro, and A. Leary. (2020). “Online display advertising markets: A literature review and future directions”. *Information Systems Research*. 31(2): 556–575.
- Choi, H. and C. F. . Mela. (2018). “Display advertising pricing in exchange markets”. *Working paper*.
- Ciocan, D. F. and V. Farias. (2012). “Model Predictive Control for Dynamic Resource Allocation”. *Mathematics of Operations Research*.
- Cole, R. and T. Roughgarden. (2014). “The sample complexity of revenue maximization”. In: *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*. 243–252.
- Conitzer, V. and T. Sandholm. (2002). “Complexity of mechanism design”. In: *Proceedings of the Eighteenth conference on Uncertainty in artificial intelligence*. 103–110.
- Cremer, J. and R. P. McLean. (1988). “Full extraction of the surplus in Bayesian and dominant strategy auctions”. In: *Econometrica: Journal of the Econometric Society*. JSTOR. 1247–1257.
- Daskalakis, C., A. Deckelbaum, and C. Tzamos. (2013). “Mechanism design via optimal transport”. In: *Proceedings of the fourteenth ACM conference on Electronic commerce*. 269–286.
- Degenne, R. and V. Perchet. (2016). “Anytime optimal algorithms in stochastic multi-armed bandits”. In: *International Conference on Machine Learning*. 1587–1595.
- Deng, Y., J. Schneider, and B. Sivan. (2019). “Prior-Free Dynamic Auctions with Low Regret Buyers”. In: *Advances in Neural Information Processing Systems*. 4804–4814.
- Devanur, N. R., Z. Huang, and C.-A. Psomas. (2016). “The sample complexity of auctions with side information”. In: *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*. 426–439.
- Devanur, N. R. and S. M. Kakade. (2009). “The Price of Truthfulness for Pay-per-Click Auctions”. In: *Proceedings of the 10th ACM Conference on Electronic Commerce. EC '09*. Stanford, California, USA: Association for Computing Machinery. 99–106.
- Dhangwatnotai, P., T. Roughgarden, and Q. Yan. (2015). “Revenue maximization with a single sample”. *Games and Economic Behavior*. 91: 318–333.

- Drutsa, A. (2020). “Reserve pricing in repeated second-price auctions with strategic bidders”. In: *International Conference on Machine Learning*. PMLR. 2678–2689.
- Dudley, R. M. (2014). *Uniform Central Limit Theorems*. Cambridge Studies in Advanced Mathematics. Cambridge University Press.
- Dütting, P., Z. Feng, H. Narasimhan, D. Parkes, and S. S. Ravindranath. (2019). “Optimal auctions through deep learning”. In: *International Conference on Machine Learning*. PMLR. 1706–1715.
- Elkind, E. (2007). “Designing and learning optimal finite support auctions”. In: *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*. 736–745.
- Epasto, A., M. Mahdian, V. Mirrokni, and S. Zuo. (2018). “Incentive-aware learning for large markets”. In: *Proceedings of the 2018 World Wide Web Conference*. 1369–1378.
- Feng, Z., S. Lahaie, J. Schneider, and J. Ye. (2021). “Reserve Price Optimization for First Price Auctions in Display Advertising”. *International Conference on Machine Learning*: 3230–3239.
- Feng, Z., H. Narasimhan, and D. C. Parkes. (2018a). “Deep learning for revenue-optimal auctions with budgets”. In: *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems*. 354–362.
- Feng, Z., C. Podimata, and V. Syrgkanis. (2018b). “Learning to bid without knowing your value”. In: *Proceedings of the 2018 ACM Conference on Economics and Computation*. 505–522.
- Fernandez-Tapia, J. (2015). “An analytical solution to the budget-pacing problem in programmatic advertising”. *Journal of Information and Optimization Sciences*. 40.
- Fernandez-Tapia, J., O. Guéant, and J.-M. Lasry. (2016). “Optimal Real-Time Bidding Strategies”. *Applied Mathematics Research eXpress*.
- Fibich, G. and A. Gavious. (2003). “Asymmetric First-Price Auctions: A Perturbation Approach”. *Mathematics of Operations Research*. 28(4): 836–852.
- Fibich, G. and N. Gavish. (2012). “Asymmetric First-Price Auctions—A Dynamical-Systems Approach”. *Mathematics of Operations Research*. 37(2): 219–243.

- Fu, H. (2013). “VCG auctions with reserve prices: Lazy or eager”. In: *Proceedings of the Fourteenth ACM Conference on Economics and Computation*.
- Fu, H. (2016). “Notes on Myerson’s Revenue Optimal Mechanisms”. URL: <http://fuhuthu.com/notes/iron.pdf>.
- Fu, H., N. Immorlica, B. Lucier, and P. Strack. (2015). “Randomization beats second price as a prior-independent auction”. In: *Proceedings of the Sixteenth ACM Conference on Economics and Computation*. 323–323.
- Gayle, W.-R. and J. F. Richard. (2008). “Numerical Solutions of Asymmetric, First-Price, Independent Private Values Auctions”. *Computational Economics*. 32(3).
- Ghosh, A., B. I. Rubinstein, S. Vassilvitskii, and M. Zinkevich. (2009). “Adaptive Bidding for Display Advertising”. In: *Proceedings of the 18th International Conference on World Wide Web. WWW '09*. 251–260.
- Golowich, N., H. Narasimhan, and D. C. Parkes. (2018). “Deep learning for multi-facility location mechanism design”. In: *Proceedings of the 27th International Joint Conference on Artificial Intelligence*. 261–267.
- Golrezaei, N., M. Lin, V. Mirrokni, and H. Nazerzadeh. (2017). “Boosted Second-price Auctions for Heterogeneous Bidders”. In: *Management Science*.
- Golrezaei, N., A. Javanmard, and V. Mirrokni. (2021). “Dynamic incentive-aware learning: Robust pricing in contextual auctions”. *Operations Research*. 69(1): 297–314.
- Gonczarowski, Y. A. and N. Nisan. (2017). “Efficient empirical revenue maximization in single-parameter auction environments”. In: *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*.
- Groeneboom, P. and G. Jongbloed. (2014). *Nonparametric Estimation under Shape Constraints*. Cambridge University Press.
- Guerre, E., I. Perrigne, and Q. Vuong. (2000). “Optimal Nonparametric Estimation of First-price Auctions”. *Econometrica*. 68(3): 525–574.
- Gummadi, R., P. Key, and A. Proutiere. (2012). *Optimal Bidding Strategies and Equilibria in Dynamic Auctions with Budget Constraints*.

- Guo, C., Z. Huang, and X. Zhang. (2019). “Settling the sample complexity of single-parameter revenue maximization”. In: *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*.
- Hartline, J., A. Johnsen, and Y. Li. (2020). “Benchmark design and prior-independent optimization”. *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*: 294–305.
- Hartline, J. D. (2013). “Bayesian mechanism design”. *Foundations and Trends® in Theoretical Computer Science*. 8(3): 143–263.
- Hartline, J. D. and T. Roughgarden. (2009). “Simple versus optimal mechanisms”. In: *Proceedings of the 10th ACM conference on Electronic commerce*. 225–234.
- Hausler, D. (1992). “Decision theoretic generalizations of the PAC model for neural net and other learning applications”. In: *Information and computation*.
- Hiriart-Urruty, J.-B. and C. Lemaréchal. (2001). *Fundamentals of Convex Analysis*.
- Huang, Z., Y. Mansour, and T. Roughgarden. (2018). “Making the most of your samples”. *SIAM Journal on Computing*. 47(3): 651–674.
- Kandasamy, K., J. E. Gonzalez, M. I. Jordan, and I. Stoica. (2020). “Mechanism Design with Bandit Feedback”. arXiv: [2004.08924](https://arxiv.org/abs/2004.08924) [stat.ML].
- Kanoria, Y. and H. Nazerzadeh. (2014). “Dynamic Reserve Prices for Repeated Auctions: Learning from Bids”. In: *Web and Internet Economics: 10th International Conference*. Vol. 8877. Springer. 232.
- Kirkegaard, R. (2009). “Asymmetric first price auctions”. *Journal of Economic Theory*. 144(4): 1617–1635.
- Kleinberg, R. and T. Leighton. (2003). “The value of knowing a demand curve: Bounds on regret for online posted-price auctions”. In: *44th Annual IEEE Symposium on Foundations of Computer Science, 2003. Proceedings*. IEEE. 594–605.
- Koltchinskii, V., D. Panchenko, *et al.* (2002). “Empirical margin distributions and bounding the generalization error of combined classifiers”. *The Annals of Statistics*. 30(1): 1–50.
- Kotowski, M. H. (2018). “On asymmetric reserve prices”. *Theoretical Economics*. 13(1): 205–237.
- Krishna, V. (2009). *Auction theory*. Academic press.



- Lattimore, T. and C. Szepesvári. (2020). *Bandit algorithms*. Cambridge University Press.
- Lavi, R. and N. Nisan. (2004). “Competitive analysis of incentive compatible on-line auctions”. *Theoretical Computer Science*. 310(1-3): 159–180.
- Le Thi, H. A., V. N. Huynh, and T. P. Dinh. (2014). “DC Programming and DCA for General DC Programs”. In: *Advanced Computational Methods for Knowledge Engineering*. Ed. by T. van Do, H. A. L. Thi, and N. T. Nguyen. Cham: Springer International Publishing. 15–35.
- Lebrun, B. (1999). “First Price Auctions in the Asymmetric N Bidder Case”. *International Economic Review*. (1).
- Lecué, G. and M. Lerasle. (2020). “Robust machine learning by median-of-means: theory and practice”. *The Annals of Statistics*. 48(2): 906–931.
- Lee, K.-C., A. Jalali, and A. Dasdan. (2013). “Real time bid optimization with smooth budget delivery in online advertising.” *Proceedings of the Seventh International Workshop on Data Mining for Online Advertising*.
- Liu, A., R. P. Leme, and J. Schneider. (2021). “Optimal contextual pricing and extensions”. In: *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM. 1059–1078.
- Lugosi, G. and S. Mendelson. (2019). “Mean estimation and regression under heavy-tailed distributions: A survey”. *Foundations of Computational Mathematics*. 19(5): 1145–1190.
- Manelli, A. M. and D. R. Vincent. (2007). “Multidimensional mechanism design: Revenue maximization and the multiple-good monopoly”. *Journal of Economic theory*. 137(1): 153–185.
- Marshall, R., M. Meurer, J. Richard, and W. Stromquist. (1994). “Numerical analysis of asymmetric first price auctions”. *Games and Economic Behavior*. (2).
- Massart, P. (1990). “The Tight Constant in the Dvoretzky-Kiefer-Wolfowitz Inequality”. In: *The Annals of Probability*. Vol. 18. No. 3. Institute of Mathematical Statistics. 1269–1283.

- Medina, A. M. and S. Vassilvitskii. (2017). “Revenue optimization with approximate bid predictions”. In: *Proceedings of the 31st International Conference on Neural Information Processing Systems*. 1856–1864.
- Milgrom, P. (2004). *Putting auction theory to work*. Cambridge University Press.
- Milgrom, P. and I. Segal. (2002). “Envelope theorems for arbitrary choice sets”. *Econometrica*. 70(2): 583–601.
- Mirroknii, V. S., R. P. Leme, P. Tang, and S. Zuo. (2016). “Dynamic Auctions with Bank Accounts.” In: *Proceedings of IJCAI*. 387–393.
- Mohri, M. and A. M. Medina. (2015). “Revenue optimization against strategic buyers”. *Advances in Neural Information Processing Systems*. 2015: 2530–2538.
- Mohri, M. and A. M. Medina. (2014). “Learning theory and algorithms for revenue optimization in second price auctions with reserve”. In: *International Conference on Machine Learning*. PMLR. 262–270.
- Morgenstern, J. and T. Roughgarden. (2015). “The pseudo-dimension of near-optimal auctions”. In: *Proceedings of the 28th International Conference on Neural Information Processing Systems-Volume 1*. 136–144.
- Myerson, R. B. (1981). “Optimal auction design”. *Mathematics of operations research*. 6(1): 58–73.
- Nazerzadeh, H., A. Saberi, and R. Vohra. (2008). “Dynamic Cost-per-Action Mechanisms and Applications to Online Advertising”. In: *Proceedings of the 17th International Conference on World Wide Web. WWW '08*. Beijing, China: Association for Computing Machinery. 179–188.
- Nedelec, T., M. Abeille, C. Calauzènes, N. E. Karoui, B. Heymann, and V. Perchet. (2019a). “Thresholding at the monopoly price: an agnostic way to improve bidding strategies in revenue-maximizing auctions”. In: *The Workshop on Learning in the Presence of Strategic Behavior, EC*.
- Nedelec, T., J. Baudet, V. Perchet, and N. E. Karoui. (2021). “Adversarial learning for revenue-maximizing auctions”. In: *20th International Conference on Autonomous Agents and Multiagent Systems*.

- Nedelec, T., N. El Karoui, and V. Perchet. (2019b). “Learning to bid in revenue-maximizing auctions”. In: *International Conference on Machine Learning*. PMLR. 4781–4789.
- Nekipelov, D., V. Syrgkanis, and E. Tardos. (2015). “Econometrics for learning agents”. In: *Proceedings of the Sixteenth ACM Conference on Economics and Computation*. 1–18.
- Ostrovsky, M. and M. Schwarz. (2011). “Reserve prices in internet advertising auctions: A field experiment”. In: *Proceedings of the 12th ACM conference on Electronic commerce*. 59–60.
- Paes Leme, R., M. Pál, and S. Vassilvitskii. (2016). “A field guide to personalized reserve prices”. In: *Proceedings of the 25th international conference on world wide web*. 1093–1102.
- Perchet, V. and P. Rigollet. (2013). “The multi-armed bandit problem with covariates”. *The Annals of Statistics*. 41(2): 693–721.
- Rahme, J., S. Jelassi, and S. M. Weinberg. (2020). “Auction learning as a two-player game”. *arXiv preprint arXiv:2006.05684*.
- Riley, J. G. and W. F. Samuelson. (1981). “Optimal auctions”. *The American Economic Review*. 71(3): 381–392.
- Rockafellar, R. T. (1970). *Convex Analysis*. Princeton University Press.
- Roughgarden, T. and O. Schrijvers. (2016). “Ironing in the dark”. In: *Proceedings of EC*. 1–18.
- Roughgarden, T. and J. R. Wang. (2016). “Minimizing Regret with Multiple Reserves”. In: *Proceedings of the 2016 ACM Conference on Economics and Computation*. 601–616.
- Rudolph, M. R., J. G. Ellis, and D. M. Blei. (2016). “Objective variables for probabilistic revenue maximization in second-price auctions with reserve”. In: *Proceedings of the 25th International Conference on World Wide Web*. 1113–1122.
- Shalev-Shwartz, S. and S. Ben-David. (2014). *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press.
- Shen, W., S. Lahaie, and R. P. Leme. (2019a). “Learning to clear the market”. In: *International Conference on Machine Learning*. PMLR. 5710–5718.

- Shen, W., P. Tang, and S. Zuo. (2019b). “Automated mechanism design via neural networks”. In: *Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems*. 215–223.
- Slivkins, A. (2019). “Introduction to Multi-Armed Bandits”. *Foundations and Trends® in Machine Learning*. 12(1-2): 1–286.
- Tang, P. and Y. Zeng. (2018). “The price of prior dependence in auctions”. In: *Proceedings of the 2018 ACM Conference on Economics and Computation*. 485–502.
- Vickrey, W. (1961). “Counterspeculation, auctions, and competitive sealed tenders”. In: *The Journal of finance*. Vol. 16. No. 1. Wiley Online Library.
- Weed, J., V. Perchet, and P. Rigollet. (2016). “Online learning in repeated auctions”. In: *Conference on Learning Theory*. PMLR. 1562–1583.
- Xu, J., K.-c. Lee, W. Li, H. Qi, and Q. Lu. (2015). “Smart pacing for effective online ad campaign optimization”. In: *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 2217–2226.
- Yao, A. C.-C. (2017). “Dominant-strategy versus bayesian multi-item auctions: Maximum revenue determination and comparison”. In: *Proceedings of the 2017 ACM Conference on Economics and Computation*. 3–20.
- Yuan, S., J. Wang, and X. Zhao. (2013). “Real-time bidding for online advertising: measurement and analysis”. *Proceedings of the Seventh International Workshop on Data Mining for Online Advertising*.