

Table of Contents

Editors' Introduction IRENE LO and SAM TAGGART	1
SIGecom Job Market Candidate Profiles 2024 VASILIS GKATZELIS and JASON HARTLINE	4
Market Design Job Market Candidate Profiles 2024 YANNAI A. GONCZAROWSKI, ASSAF ROMM, and RAN SHORRER	36
SIGecom Winter Meeting 2024 Highlights AGHAHEYBAT MAMMADOV, EMILY RYU, and ROBERTO SAITTO	42
Decision Theory, Subjective Uncertainty, and Computer Science GIACOMO LANZANI	59
Welfare-Maximizing Pooled Testing SIMON FINSTER, MICHELLE GONZÁLEZ AMADOR, EDWIN LOCK, FRANCISCO MARMOLEJO-COSSÍO, EVI MICHA, and ARIEL D. PROCACCIA	66
Leveraging Reviews: Learning to Price with Buyer and Seller Uncertainty WENSHUO GUO, NIKA HAGHTALAB, KIRTHEVASAN KANDASAMY, and ELLEN VITERCIK	74
Inequality and Market Design PIOTR DWORCZAK	83
Generative AI as Economic Agents NICOLE IMMORLICA, BRENDAN LUCIER, and ALEKSANDRS SLIVKINS	93
Causal Inference under Incentives: An Annotated Reading List KEEGAN HARRIS and VASILIS SYRGKANIS	110
Impartial Peer Selection: An Annotated Reading List OMER LEV, HARPER LYON, and NICHOLAS MATTEI	113
Assortment Optimization: An Annotated Reading Assortment WILL MA	118
Recent Trends in Information Elicitation RAFAEL FRONGILLO and BO WAGGONER	122
Online Matching: A Brief Survey ZHIYI HUANG, ZHIHAO GAVIN TANG, and DAVID WAJC	135

Auto-bidding and Auctions in Online Advertising: A Survey	159
GAGAN AGGARWAL, ASHWINKUMAR BADANIDIYURU, SANTIAGO R. BALSEIRO, KSHIPRA BHAWALKAR, YUAN DENG, ZHE FENG, GAGAN GOEL, CHRISTOPHER LIAW, HAIHAO LU, MOHAMMAD MAHDIAN, JIEMING MAO, ARANYAK MEHTA, VAHAB MIRROKNI, RENATO PAES LEME, ANDRES PERLROTH, GEORGIOS PILIOURAS, JON SCHNEIDER, ARIEL SCHVARTZMAN, BALASUBRAMANIAN SIVAN, KELLY SPENDLOVE, YIFENG TENG, DI WANG, HANRUI ZHANG, MINGFEI ZHAO, WENNAN ZHU, and SONG ZUO	

ACM SIGecom Exchanges, Vol. 22, No. 1, June 2024

Editors-in-Chief: Irene Lo and Sam Taggart

Information Director: Yannai A. Gonczarowski

ACM Staff: Irene Frawley

Notice to Contributing Authors to SIG Newsletters

As a contributing author, you retain copyright to your article. ACM will refer all requests for republication directly to you.

By submitting your article for distribution in any newsletter of the ACM Special Interest Groups, you hereby grant to ACM the following non-exclusive, perpetual, worldwide rights:

- to publish your work online or in print on condition of acceptance by the editor
- to include the article in the ACM Digital Library and in any Digital Library-related services
- to allow users to make a personal copy of the article for noncommercial, educational, or research purposes
- to upload your video and other supplemental material to the ACM Digital Library, the ACM YouTube channel, and the SIG newsletter site

Furthermore, you affirm that:

- if third-party materials were used in your published work, supplemental material, or video, that you have the necessary permissions to use those third-party materials in your work

Editors' Introduction

IRENE LO

Stanford University

and

SAM TAGGART

Oberlin College

This issue of SIGecom Exchanges, we are excited to present an expanded edition which combines the Fall 2023 and Spring 2024 issues. On the news side, this issue celebrates the 2024 slate of job candidates with the annual list of candidate profiles. This is followed by a report on the SIGecom 2024 Winter Meeting, which includes talk summaries and interviews with invited speakers. On the research side, the issue includes five research letters, three annotated reading lists, and three surveys.

The 2024 SIGecom Job Candidates Profiles are compiled by Vasilis Gkatzelis and Jason Hartline. This is the ninth year of this wonderful annual tradition. In addition, we have the 2024 Market Design Job Candidate Profiles, compiled by Yannai Gonczarowski, Assaf Romm, and Ran Shorrer. This is the fourth annual collection from this community. Both sets of profiles have been available since Fall 2023, and for completeness we include them in this expanded Spring 2024 issue. Many thanks to Vasilis, Jason, Yannai, Assaf, and Ran for their service to the community in creating these collections each year.

A highlight of the issue is the report on the fourth SIGecom Winter Meeting, written by graduate students Aghaheybat Mammadov, Emily Ryu, and Roberto Saitto. The event took place virtually in February 2024 and was on the topic of behavioral models. The report features an excellent collection of event summaries and speaker interviews. The summaries are comprehensive and cover all the events at the meeting: introductory talks by Jon Kleinberg and Ori Heffetz, a fireside chat with Noam Nisan and Al Roth, and spotlight talks by Modibo Camara, Ryan Oprea, Gali Noti, and Nicole Immorlica. Interviews with Noam Nisan, Modibo Camara, and Nicole Immorlica provide additional insight on what inspired the research ideas presented in the talks, promising areas for future research, as well as lesser-known fun facts about each interviewee.

This issue of SIGecom exchanges includes five research letters, which highlight several of the award-winning papers from EC 2023, as well as exciting emerging areas of research.

—Giacomo Lanzani, winner of the 2023 EC Best Paper with a Student Lead Author Award, provides an intriguing letter discussing connections between Decision Theory and Computer Science. He uses his EC'24 paper as an illustration of how questions in the former can be addressed using methods from the latter: specifically, how evaluation criteria from computer science can help characterize appropriate dynamic decision criteria for an agent who wants a decision rule that

Author's address: ilo@stanford.edu, staggart@oberlin.edu.

is robust to uncertainty and model misspecification.

- A letter from Simon Finster, Michelle González Amador, Edwin Lock Francisco Marmolejo Cossio, Evi Micha, and Ariel Procaccia summarizes their paper which was awarded the 2023 EC Exemplary Applied Modeling Track Award. This paper studies the optimal use of pooled testing for diverse populations during pandemics, considering individual infection probabilities and utility for in-person activities. The authors find that non-overlapping testing allocations, which are simpler to implement, are near-optimal and effective in practice. They also successfully pilot a practical application of utility-weighted testing in a real-world setting.
- Wenshuo Guo, Nika Haghtalab, Kirthivasan Kandasamy, and Ellen Vitercik provide an overview of their paper which received the 2023 EC Exemplary Artificial Intelligence Track Award. The paper considers algorithmic learning in a pricing problem where buyers depend on reviews for information about the product. The introduction of reviews adds a new dynamic to the seller’s learning problem: in addition to the usual explore-exploit tradeoff inherent in online pricing, the algorithm now needs to ensure that there are sufficient reviews to allow buyers to infer their values. Only with enough reviews can the seller accurately infer demand and set prices well. The paper gives algorithms with tight theoretical regret bounds.
- A research letter from Piotr Dworczak considers the problem of inequality-aware market design. Drawing on two of his recent papers, he provides a simple framework for designing regulation and other policies to ensure equitable redistribution in markets with asymmetric or imperfect information. He also presents multiple directions for future research that will be of broad interest in the EC community.
- Nicole Immorlica, Brendan Lucier, and Alex Slivkins conclude the letters section with a position piece on generative AI. They observe that a main application for AI technologies is as *consultant* or *agent*, often reducing the costs of information acquisition or making certain tasks easier. At the same time, though, these technologies need not have objectives that are directly aligned with the decision-maker. The authors argue that these applications are excellent opportunities to apply and further develop our understanding of delegation and related topics in contract theory. The piece suggests fascinating avenues for future research.

For readers looking for directed summer reading, this issue also includes annotated reading lists on three great topics.

- A reading list from Keegan Harris and Vasilis Syrgkanis provides an introduction to the growing area of causal inference in the presence of strategic agents. The reading list includes seminal papers from statistics and econometrics, as well as recent papers bringing in tools from machine learning and game theory.
- Omer Lev, Harper Lyon, and Nicholas Mattei give a guided tour of the literature on peer selection. This is an area of mechanism design without money that has seen a recent burst of interest. A group of agents must select one or more of their own to serve in an elected role or receive an award. The goal is to design mechanisms that select a worthy candidate in a non-manipulable way.

- An annotated reading assortment from Will Ma serves as an excellent primer for members of our community interested in the problem of assortment optimization. The goal is to identify assortments of products to show customers (whose choices are given by parametric choice functions) to maximize expected revenue. The assortment includes seminal papers, new choice models, online variants of the problem, as well as future directions that may be of interest to our readers.

Finally, the issue includes three excellent research surveys.

- Bo Waggoner and Raf Frongillo provide a short but highly informative overview of the area of information elicitation. The aim is to design scoring and reward functions to incentivize accurate reporting of statistical properties by an informed agent. For both single- and multi-agent problems, they highlight both unsolved technical questions and connections to other fields of interest.
- Zhiyi Huang, David Wajc, and Zhihao Gavin Tang give a thorough survey of the very deep literature on online matching. They walk us carefully through a wide range of variants, and highlight the best-known algorithms for each. The survey includes an appendix with a “teachable moment:” a lesson-length analysis of the well-studied Balance algorithm for fractional online matching.
- Members of the EC community working at Google provide a comprehensive survey of auto-bidding and auctions in online advertising. They define the problem space, and identify the main challenges faced by the bidding agent and auctioneer. They then survey recent literature addressing each of these challenges: for agents, how to bid optimally in various settings; and for auctioneers, equilibrium outcomes and auction design. The survey closes with explorations of multiple emerging topics, including determining the utility functions of autobidders, identifying the effects of multiple advertising channels, and findings from empirical studies of different auction formats.

Many thanks to Yannai Gonczarowski for his help in putting together this issue, and thanks to our contributors for their efforts and expertise. As always, please continue to volunteer letters, surveys, annotated reading lists or position papers. We hope you enjoy this expanded issue and find inspiration in the rich and varied research areas being explored by members of our community!

Yurong Chen private information games, Stackelberg equilibria, online ad auctions	11
Michael Curry machine learning, mechanism design, differentiable economics, auctions	12
Sulagna Dasgupta mechanism design, test design, strategic communication, hard evidence	13
Tom Demeulemeester operations research, randomization, matching	14
Zhun Deng responsible ML, computational social science, statistics	15
Kate Donahue algorithmic fairness, game theory, machine learning	16
Jessie Finocchiaro surrogate loss functions, algorithmic decision-making, consistency	17
Bailey Flanigan social choice, algorithmic game theory, deliberation, algorithms	18
Yuri Resende Fonseca digital platforms, revealed preferences, online learning	19
Abheek Ghosh contests, dynamics in games, mechanism design, equilibrium complexity	20
Denizalp (Deni) Goktas multiagent learning, equilibrium computation, general equilibrium theory	21
Yingkai Li mechanism design, information design, online algorithms	22
Zun Li equilibrium computation, multi-agent reinforcement learning	23
Francisco Marmolejo-Cossío resource allocation, blockchain, mechanism design for social good	24
Vishnu V. Narayan fair division, auctions, online algorithms, discrete optimization	25
Orestis Papadigenopoulos online decision-making and learning, revenue management, algorithms	26

Justin Payan	
market design, fair division, social good, natural language processing	27
Farzad Pourbabaee	
mechanism design, information economics, game theory, statistics	28
Rojin Rezvan	
mechanism design, fair auction design, game theory	29
Abhin Shah	
causal inference, algorithmic fairness, personalized decision-making	30
Kangning Wang	
mechanism design, social choice, information design, algorithmic fairness	31
Lily Xu	
ML, multi-armed bandits, reinforcement learning, AI for social impact	32
Aviv Yaish	
blockchain, security, mechanism design, decentralized finance	33

YEGANEH ALIMOHAMMADI ([Homepage](#), [CV](#))

Thesis: Learning and Decision Making using Network Data ('24)

Advisor: Amin Saberi, Stanford University

Brief Biography: Yeganeh is a final-year Ph.D. candidate in Stanford University's Operations Research group. Before joining Stanford, she obtained a B.Sc. in Computer Engineering with a Mathematics minor from Sharif University of Technology. In 2022, she was a research fellow at the Simons Institute and interned at Google Research, hosted by Aranyak Mehta.

Research Summary: Yeganeh's research delves into the analysis of large-scale networks and stochastic systems. She employs tools from applied probability and algorithm design to tackle pressing challenges in business operations.

In today's interconnected world, networks from supply chains to social media play crucial roles. Yeganeh explores how *micro-level* behaviors in these networks cascade into *macro-level impacts*. She has provided insights into how localized data from a handful of individuals can predict epidemics [1,2]; how individual advertising strategies shape system-wide outcomes in online auctions [3]; and how adding a few drivers influences overall efficiency in ridesharing platforms [4].

Another dimension of her research is on creating *sampling algorithms* to streamline the analysis of expansive networks. Yeganeh designs these algorithms to extract a representative subset of nodes and edges that maintain the key topological characteristics of the underlying network and enable fast computation. Beyond using these for epidemic forecasting [1], she has formulated sampling algorithms for enhancing graph neural networks and sampling specific network configurations [5].

Finally, a critical component of her research involves translating theoretical insights into actionable strategies. Using network models, she developed prediction methods that informed the COVID-19 reopening strategies of Los Angeles Unified School District (LAUSD) – the nation's second-largest school district [6].

Representative Papers:

- [1] Epidemic Forecasting on Networks: Bridging Local Samples with Global Outcomes (submitted to Operations Research)
with C. Borgs, R. van der Hofstad, and A. Saberi
- [2] Locality of Random Digraphs on Expanders (The Annals of Probability, 2023)
with C. Borgs and A. Saberi
- [3] Incentive Compatibility in the Auto-Bidding World (EC, 2023)
with A. Mehta and A. Perloth
- [4] The Value of Excess Supply in Spatial Matching Markets (EC, 2022)
with M. Akbarpour, S. Li, and A. Saberi
- [5] Sequential Importance Sampling for Estimating Expectations over the Space of Perfect Matchings (The Annals of Applied Probability, 2023)
with P. Diaconis, M. Roghani, and A. Saberi
- [6] Relative-Risk and the Assessment of School Safety in the COVID-19 Pandemic (Health Management, Policy, and Innovation, 2021)
with K. Shiragur, R. Johari, K. Schulman, and K. Staudenmayer

HEDYEH BEYHAGHI ([Homepage](#), [CV](#), [Google Scholar](#))

Thesis: Approximately-Optimal Mechanisms in Auction Design, Search Theory, and Matching Markets ('19)

Advisor: Éva Tardos

Brief Biography: Hedyeh Beyhaghi is a postdoctoral research associate in the School of Computer Science at Carnegie Mellon University, hosted by Nina Balcan. Hedyeh received her PhD in Computer Science from Cornell University, advised by Éva Tardos. During her PhD studies, she was a long-term visitor at the Simons Institute for the Theory of Computing for the Economics and Computation program, an intern at Google, and an Ivy-Plus Exchange Scholar at Princeton University hosted by Matt Weinberg. Before joining CMU, Hedyeh was a postdoctoral research fellow at Toyota Technological Institute at Chicago (TTIC) and Northwestern University, hosted by Avrim Blum, Jason Hartline, and Samir Khuller.

Research Summary: My research primarily revolves around decision-making in the context of strategic behavior and uncertainty. This involves creating solutions when the input is derived from strategic agents, devising rules to counteract self-interested agents, and developing algorithms with limited knowledge about future events or based on limited data. Such problems are prevalent in numerous real-world situations, ranging from learning classification rules and designing auctions to online stochastic optimization.

Decision-making under strategic behavior and uncertainty poses several challenges. The solutions must ensure the agent's actions align with the overall good outcome, even when they are driven by their own self-interest. Agents may manipulate their characteristics to achieve a better outcome, and there may not be access to any unmanipulated data. The agents' behavior can vary in response to decision-making parameters. Furthermore, while the decisions made are often irreversible in a sequential scenario, the solutions need to be robust and resilient, able to perform well in unforeseen future scenarios.

To address these complex issues, I utilize a combination of techniques from mechanism design, online optimization, and learning theory. By leveraging these approaches, I design efficient algorithms that perform competently in uncertain and strategic environments. My contributions include: (1) Providing novel characterizations and solutions to online optimization problems under uncertainty that relax assumptions for more realistic, more general scenarios. (2) Successful adaptation of classic learning algorithms for strategic scenarios, including the first high-performance learning algorithms in several settings. (3) Designing algorithms with enhanced societal aspects, e.g., improving social development and fairness, in addition to traditional optimization criteria.

Representative Papers:

- [1] Pandora's Problem with Nonobligatory Inspection: Optimal Structure and a PTAS (STOC 2023) with L. Cai
- [2] The Strategic Perceptron (EC 2021) with S. Ahmadi, A. Blum, and K. Naggita
- [3] Improved Revenue Bounds for Posted-Price and Second-Price Mechanisms (Operations Research 2021) with N. Golrezaei, R. Paes Leme, M. Pal, and B. Sivan

JOHANNES BRÜSTLE ([CV](#), [Google Scholar](#))

Thesis: The Competition Complexity of Online Mechanisms ('24)

Advisors: Paul Dütting (Google Research Zurich), László Végő (London School of Economics and Political Science)

Brief Biography: I am currently in my fourth year of pursuing a Ph.D. at the London School of Economics and Political Science, under the guidance of Professors Paul Dütting and László Végő. My research is supported by an LSE PhD Studentship.

Research Summary: I am broadly interested in algorithmic game theory, mechanism design and the asymptotics of online decision making. During my PhD, I have been investigating the prophet inequality setting with additional bidders for the online algorithm. Furthermore, I have been lucky to be able to work on different subjects within algorithmic game theory, such as robust mechanisms and their connection to learning as well as fair division of indivisible goods.

In [2], we study the problem of comparing the performance of the optimal posted price mechanism for single item online arrival to that of the prophet. We explore a new direction compared to the well known results in this area by giving the online algorithm an additional number of i.i.d bidders. How many such additional resources are necessary for the online algorithm to catch up to the performance of the prophet? Surprisingly, we find that we can analyze the corresponding optimization problem exactly. In a working paper [1] we find that by a very different approach, we recover the same asymptotic behavior even for the independent non-i.i.d setting.

I am also interested in learning within algorithmic game theory. More precisely, consider the classical problem of helping an auctioneer obtain large expected revenue through an auction in which the bidders have incentive to participate and to report their values truthfully for any bundles of items offered. In [3], we show that robust mechanisms allow us to push the boundary of learning (approximately) optimal multi-item auctions to the important setting of item dependence.

Another area I have worked on is fair division of indivisible goods. In [4], we want to get an upper bound on the total amount of money that has to be given to agents on top of their allocations in order to eliminate envy. We achieve the tight upper bound on total subsidy which is $(n - 1)$ dollars, where the most valued item is normalized to 1 dollar.

Representative Papers:

- [1] The Competition Complexity of Prophet Inequalities (working paper)
with J. Correa, P. Dütting, T. Ezra, M. Feldman and V. Verdugo
- [2] The Competition Complexity of Dynamic Pricing (EC 2022, MOR 2023)
with J. Correa, P. Dütting, and V. Verdugo
- [3] Multi-Item Mechanisms without Item-Independence: Learnability via Robustness (EC 2020) with Y. Cai, and C. Daskalakis
- [4] One Dollar Each Eliminates Envy (EC 2020)
with J. Dippel, V.V. Narayan, M. Suzuki, and A. Vetta

LINDA CAI ([Homepage](#), [CV](#), [Google Scholar](#))

Thesis: Strategic Decision Making in the Wild: Information Acquisition, Resource Augmentation and Practical Irrationality ('24)

Advisor: Matt Weinberg, Princeton University

Brief Biography: I am currently a fourth (and final) year PhD student in Computer Science at Princeton University advised by Prof. Matt Weinberg. My research interest includes a broad set of problems that relate to online optimization, incentive, learning, and their intersection. I am fortunate to be the recipient of a number of scholarships and awards, including School of Engineering and Applied Science Award for Excellence, Siebel Scholar, Chainlink Fellowship and Francis Robbins Upton Fellowship.

Research Summary: Recently, my work focuses on understanding the impact of relaxing or changing key assumptions in economic models. For instance, I have studied the effect of resource augmentation, existence of non-rational (possibly learning) agents, and the role of inspection in optimal search setting.

In [1], we study the classical setting of optimal search amidst search friction, but with the tweak that pre-selection information acquisition isn't mandatory. Our results reveal a structured optimal policy that requires minimal order adaptation, and where computational hardness arises primarily from determining the initial inspection order. We also resolve a longstanding open problem regarding the approximability of optimal utility in this setting by providing a polynomial approximation scheme (PTAS).

In the repeated mechanism design setting, my interest lies in scenarios involving not fully rational participants, such as those using learning algorithms from past interactions, and their susceptibility or resistance to exploitation by mechanism designers. For instance, in [2], we show that auctioneers can extract the maximum possible revenue (close to the optimal social welfare) against multiple unsophisticated no-regret learners, but face complexities and reduced revenue extraction when learners cap their bids to item values.

I have also investigated the impact of resource augmentation in the traditional one-shot auction design setting. While a complex auction may extract more revenue from a fixed bidder pool, its complexity might dissuade participation. The auctioneer hence need to consider the trade-off between simplicity and optimality. Our work in [3] shows that any constant revenue competitive simple auction, when scaled with a constant multiplicative increase in bidders, surpasses the original optimal revenue, except when another simple auction (second-price auction) already achieves near-optimal revenue.

Representative Papers:

- [1] Pandora's Problem with Nonobligatory Inspection: Optimal Structure and a PTAS. (STOC 2023) with Hedyeh Beyhaghi
- [2] Selling to Multiple No-Regret Buyers (WINE 2023)
with S. Matthew Weinberg, Evan Wildenhain and Shirley Zhang
- [3] 99% Revenue with Constant Enhanced Competition. (EC 2021)
with Raghuvansh R. Saxena

YURONG CHEN ([Homepage](#), [CV](#), [Google Scholar](#))

Thesis: Games of Private Information over Learning Agents ('24)

Advisor: Xiaotie Deng, Peking University

Brief Biography: Yurong is a final-year Ph.D. candidate at Peking University, advised by Xiaotie Deng. During grad school, she visited the University of Hong Kong in Spring 2023, hosted by Zhiyi Huang. Before that, she completed her B.S. degree in mathematics in Hua Luogeng Honors Class at Beihang University, where she achieved top rank in class. Yurong is a recipient of the Best Student Paper Award at WINE 2022.

Research Summary: Yurong mainly focused on understanding, learning, and computing how agents make use of their information advantage during interaction with learning agents, from a game theorist's point of view. For example, machine learning enables the extraction of real-world game parameters from data, e.g., strategy sets and payoff functions. This makes game theory possible to be applied to real-world scenarios. However, agents with information advantage can provide fake data, manipulating the learning results and game outcomes to benefit themselves. Players are already gaming with their private information during the learning process before the learned game starts. It even remains a question whether machine learning can produce credible results under such manipulation. Therefore, studying how agents utilize their information advantage is crucial for both learning and game theory to be more robustly applied to real-world situations.

Private Information Misreporting in Stackelberg equilibria [1,3]. [1] studies how the follower optimally misreports his payoff function in extensive-form games. In four settings considered, they fully characterize all the game outcomes inducible through the follower's misreporting payoffs, and provide polynomial-time algorithms to find the optimal way of manipulation. [3] studies the query version of this problem in bimatrix games: the follower does not know the leader's payoff function, but has to learn to misreport optimally. They show that a polynomial number of queries and operations is sufficient for optimally misreporting.

Coordinated Auto-Bidding in Online Ad Auctions [2]. In online ad auctions, advertisers are delegating bidding tasks to bidding agencies. While most studies on online bidding focus on bidding for one bidder in one campaign, the bidding agency actually has information on all her clients and can coordinate bids to benefit everyone. [2] proposes coordinated online bidding strategies in repeated second-price auctions with budgets that theoretically and experimentally guarantee everyone a better utility than the best they get under independent bidding.

Representative Papers:

- [1] Optimal Private Payoff Manipulation against Commitment in Extensive-form Games (WINE 2022, Best Student Paper)
with X. Deng, and Y. Li
- [2] Coordinated Dynamic Bidding in Repeated Second-Price Auctions with Budgets (ICML 2023)
with Q. Wang, Z. Duan, H. Sun, Z. Chen, X. Yan and X. Deng
- [3] Learning to Manipulate a Commitment Optimizer (Working Paper)
with X. Deng, J. Gan, and Y. Li

MICHAEL CURRY ([Homepage](#), [CV](#), [Google Scholar](#))

Thesis: Learning and Robustness With Applications to Mechanism Design ('22)

Advisor: John Dickerson & Tom Goldstein, University of Maryland

Brief Biography: Michael Curry is a postdoc splitting time between Sven Seuken's group at the University of Zurich and David Parkes' group at Harvard. His research is at the intersection of machine learning and mechanism design, with a particular focus on using machine learning techniques to search through complex spaces of mechanisms and strategies.

Research Summary:

It's long been observed that mechanism design given samples from the valuation distribution is in essence a learning problem. And mechanisms are just functions, so why not represent them using the computational tools and rich function approximators that modern deep learning provides? This is the pitch of differentiable economics, a recent thread of work introduced in a number of papers, notably "Optimal Auctions Through Deep Learning" by Dütting et al., which trains mechanisms represented directly as deep neural networks.

One thread of my research on differentiable economics has involved overcoming a significant limitation of these approaches – they are not perfectly strategyproof. In our paper "Certifying Strategyproof Auction Networks", we adapt neural auction architectures to make them certifiable – we can *exactly* compute the degree to which strategyproofness is really violated. In our IJCAI-2023 paper "Differentiable Economics for Randomized Affine Maximizer Auctions", we present a new and high-performing architecture that is also perfectly strategyproof for any setting of the parameters. (This paper has itself been further improved upon by a separate group from ours in "A Scalable Neural Network for DSIC Affine Maximizer Auction Design" by Duan et al., appearing this year at NeurIPS.)

Other work with collaborators has included devising a practical method for identifying only the most important transplants for pre-screening in paired kidney exchange. Additionally, while interning at Salesforce Research, I worked with the AI Economist team on a project to use GPU-accelerated simulations to model a simple economy involving hundreds of learning agents.

In ongoing work at Sven Seuken's group, we aim to find high-performing *dynamic* mechanisms in very general MDP settings; in a separate project, we also use RL to find good bidding strategies in iterative combinatorial auctions. And in ongoing work with David Parkes and Zhou Fan, we study a particular case of multi-good automated market making under adverse selection. We have shown that in such settings, the problem of finding a profit-maximizing automated market maker is dual to an optimal transport problem, and that differentiable economics can be used to search for optimal mechanisms.

Representative Papers:

- [1] Differentiable Economics for Randomized Affine Maximizer Auctions (IJCAI '23) with M. Curry, J. Dickerson, T. Sandholm
- [2] Certifying Strategyproof Auction Networks (NeurIPS '20) with M. Curry, P.Y. Chiang, T. Goldstein, and J. Dickerson

SULAGNA DASGUPTA ([Homepage](#), [CV](#))

Thesis: Screening Knowledge ('24)

Advisor: Ben Brooks, University of Chicago

Brief Biography: Sulagna Dasgupta is a final year PhD candidate at the University of Chicago. She is a microeconomic theorist with interests in mechanism design and information economics. She has three broad, interconnected research agendas – (A) screening knowledge, (B) the interaction of hard and soft information in strategic communication and information design, and (C) matching theory.

Research Summary: Tests of knowledge are ubiquitous – job interviews, exams, standardized tests etc. The goal of such tests is to discern some underlying quality of the test-taker, which enables them to do well on the test. My job market paper asks: How to optimally design such tests to maximize this learning? To answer it, I set up a model of screening knowledge. In order to focus on the forces of incentive new to this problem, I use its most elemental version, where the subject matter is simply a binary fact. I show that even though the test-designer can set up any complicated grading scheme by rewarding test-takers according to the strength of their conviction in the correct answer, optimally, the tests take the simple True/False or True/False/“I don’t know” forms.

The optimal tests feature two notable features. First, they may reward the admission of ignorance, by way of the “I don’t know” option. This result can be interpreted as providing a basis for the “guessing penalty” used in many multiple-choice tests such as the SAT (till 2016). Secondly, they may pass incorrect answers or fail correct answers. This “unfairness” of the optimal tests is the most salient consequence of agency issues particular to the knowledge screening problem. I show that only “obvious” answers can be failed in spite of being correct and only “counterintuitive” answers can be passed in spite of being incorrect. An obvious (respectively, counterintuitive) answer is the a priori likely (respectively, unlikely) answer, when the prior is sufficiently extreme. This reflects the common feature of real world evaluation schemes which sometimes attach greater penalty to getting “obvious” questions wrong than to getting “trick” questions wrong.

In two follow up papers (works in progress) I explore a case where (1) the test-taker also needs to “show his work” and (2) another, where the subject matter is more complex – namely, a finitely-valued state. I model the first as a mechanism design with evidence problem. Using a symmetric setting, I show that the optimal test passes (respectively, fails) the test-taker regardless of his answer, if the amount of evidence provided (“steps shown”) is sufficiently high (respectively, sufficiently low), and passes if and only if his answer is correct, when the amount of evidence provided is intermediate. In the second, again using a symmetric setting, I show that when there are potentially many questions and/or the answer possibilities of each question are many, the optimal test allows the test-taker to pick a fixed number of “correct” answers, and passes him if the actual correct answer is one of them.

Representative Papers:

- [1] Optimal Test Design for Knowledge-based Screening (EC '23)
- [2] Communication via Hard and Soft Information (EC '23)
- [3] Information Design in One-sided Matching Markets (R&R, JME)

TOM DEMEULEMEESTER ([Homepage](#), [CV](#), [Google Scholar](#))

Thesis: Fairness Through Randomization: an Operations Research Perspective ('24)

Advisor: Roel Leus (KU Leuven), Dries Goossens (Ghent University)

Brief Biography: I am a final year Ph.D. student at KU Leuven. While my educational background lies in operations research, I enjoy interdisciplinary research on the intersections with theoretical economics and theoretical computer science.

Research Summary: The main focus of my research has been to study settings where fairness can only be obtained by introducing randomization. More specifically, I have studied how to obtain probabilistic solution concepts with compelling fairness properties, and how these probabilistic outcomes can be written as a lottery over deterministic outcomes that all satisfy a range of properties that are deemed desirable for the problem at hand (“best of both worlds” approach).

A first problem that I have been particularly interested in is the one-sided matching problem (school choice, house allocation. . .). In [1], for example, we propose an algorithmic framework to decompose a given probabilistic assignment over deterministic matchings in such a way that the risk of ending up with an “undesirable” outcome is minimized, for various interpretations of undesirability. Alternatively, in [2], we introduce a new probabilistic mechanism for the classical assignment problem which is concerned with the egalitarianism of the final assignment. The resulting *Rawlsian assignment* lexicographically maximizes the well-being of the worst-off agents. We study its axiomatic properties, and validate its performance on data from housing cooperatives in Uruguay.

Second, we study how to fairly select one of the (possibly many) optimal solutions of an integer programming formulation in [3]. While solvers traditionally return one of the optimal solutions deterministically, we argue why a randomized approach is desirable for high-impact problems, such as kidney exchange, matching variants, or sortition. We propose solution methods to implement, for example, the Nash maximum welfare, the leximin or the Random Serial Dictatorship rules without requiring a full enumeration of the optimal solutions.

In general, I aim to obtain a holistic understanding of the problems that I study by collaborating with researchers from various backgrounds. This has resulted, for example, in collaborations on coalition formation [4] and committee voting [5].

Representative Papers:

- [1] A Pessimist’s Approach to One-Sided Matching (European Journal of Operational Research) with D. Goossens, B. Hermans, and R. Leus
- [2] Rawlsian Assignments (under review) with J. S. Pereyra
- [3] Fair Integer Programming Under Dichotomous Preferences (major revision, European Journal of Operational Research) with D. Goossens, B. Hermans, and R. Leus
- [4] Relaxed Core Stability in Hedonic Games with Size-Dependent Utilities (MFCS 2023) with J. Peters
- [5] Strategy-Proofness and Proportionality in Party-Approval Multi-Winner Elections (AAAI 2023) with T. Delemazure, M. Eberl, J. Israel, and P. Lederer

ZHUN DENG ([Homepage](#), [CV](#))

Thesis: Robustness, Generalization and Fairness in Learning: Analysis and Design ('22)

Advisor: Cynthia Dwork, Harvard University

Brief Biography: Zhun Deng is a postdoctoral researcher at Columbia University, and also part of Simons Collaboration on the Theory of Algorithmic Fairness. Previously, he completed his Ph.D. in the Theory of Computation group at Harvard University, advised by Cynthia Dwork. He is also fortunate to work with David Parkes, Weijie Su, and James Zou on various projects. His papers have won multiple honors such as Spotlight and Oral Presentation at flagship machine learning conferences, including ICML, NeurIPS, ICLR, and AISTATS.

Research Summary: Modern digital systems powered by artificial intelligence (AI) are facing pivotal challenges standing in the way of reliable deployment: (1) Effectively *quantifying the quality of predictions* and *enabling the control of catastrophic outcomes*, as AI is deployed in high-impact, risk-sensitive domains like medicine and autonomous driving. (2) Reducing the potential for *negative social impact* of AI systems, as they inadvertently amplify existing social biases or blindly push users towards unhealthy consumption on social platforms.

My research aims to develop the next generation of principled methods that guide the design and deployment of AI systems in a responsible and societally beneficial way. In pursuit of this goal, my research agenda centers around introducing *rigorous statistical tools* in *uncertainty estimation* and *stochastic decision-making* to ensure reliable *real-world practice* of AI systems regarding prediction, decision-making, and societal value alignment. I am particularly keen on applying these methods to *interdisciplinary domains* such as healthcare, digital sociology, and large language models.

Uncertainty quantification for reliable decision-making. I work on advancing the theoretical understanding of model uncertainty [1] and deriving novel methods [2] to endow decision-making systems with performance and uncertainty guarantees (e.g., for important risk measures like value-at-risk (VaR)/ conditional value-at-risk (CVaR)), especially for complex black-box models' (e.g., large language models) deployment.

Building socially responsible digital systems. I also work on identifying ethical and societal concerns in deploying AI in digital systems and develop algorithmic solutions to provably mitigate such concerns (see [3]).

Representative Papers:

- [1] Happymap: A Generalized Multi-Calibration Methods (ITCS'23) with C. Dwork, and L. Zhang
- [2] Distribution-Free Statistical Dispersion Control for Societal Applications (NeurIPS'23, Spotlight) with T. Zollo, J. Snell, T. Pitassi, and R. Zemel
- [3] FIFA: Making Fairness More Generalizable in Classifiers Trained on Imbalanced Data (ICLR'23) with J. Zhang, L. Zhang, T. Ye, Y. Coley, W. Su, and J. Zou

KATE DONAHUE ([Homepage](#), [CV](#), [Google Scholar](#))

Thesis: AI as a Resource: Strategy, Uncertainty, and Societal Welfare ('24)

Advisor: Jon Kleinberg, Cornell University

Brief Biography: Kate Donahue is a final year computer science PhD candidate at Cornell advised by Jon Kleinberg. She works on algorithmic problems relating to the societal impact of AI such as fairness, human/AI collaboration and game-theoretic models of federated learning. Her work has been supported by an NSF fellowship and recognized by a FAccT Best Paper award. During her PhD, she has interned at Amazon, Google, and Microsoft Research.

Research Summary: Artificial intelligence (AI) is emerging as an important new step in a long lineage of innovations transforming human society - with effects that can be positive, as well as negative. My research agenda is to study the societal impact of AI, viewing viewing AI as a resource that can be used for social welfare, among multiple self-interested agents operating under uncertainty. My goal is to provide clear, generalizable insights for how to best use AI/algorithms in a wide range of situations.

In the first vein of research, I study data as a resource, through the lens of “model-sharing games”. In this setting, each agent has access to data that can be used to create an ML model (e.g. each agent may be a hospital with data on patient outcomes). The agent can choose to build the model using only their data, or collaborate with other hospitals with related data, reflecting scenarios like federated learning and data cooperatives. However, there may be data shift between agents (e.g. different patient distributions), which means the same model performs differently for each agent, and collaborating may not always be optimal. In [KD AAAI '21], we propose a theoretical framework for studying this scenario, called “model-sharing games”, which we then build upon to explore questions around stability, optimality, fairness, and incentive-compatibility [1, 2].

Beyond the model itself, I also consider the broader system the AI is embedded in: specifically, models as inputs to human decisions. Here, we view human time as a resource and try to maximize the performance of the combined human-AI system [3]. One natural way for two agents to interact is filtering: one agent (i.e algorithm) narrows the items into a subset, from which the second agent (i.e. human) selects the final item, a process which occurs in settings as varied as product recommendation, route selection, or categorical prediction. In ongoing work with K. Kollias and S. Gollapudi, we show how performance of an AI-human system is affected by accuracy gaps between the human and algorithm, the size of the set the algorithm presents to the human, and common human cognitive biases.

Representative Papers:

- [1] Optimality and Stability in Federated Learning (Neurips 2021)
with J. Kleinberg
- [2] Models of fairness in federated learning (WWW '23)
with J. Kleinberg
- [3] Human-Algorithm Collaboration: Achieving Complementarity and Avoiding Unfairness (FAccT '22)
with A. Chouldechova, K. Kenthapadi

JESSIE FINOCCHIARO ([Homepage](#), [CV](#), [Google Scholar](#))

Thesis: Designing Consistent and Convex Surrogates for General Prediction Tasks ('22)

Advisor: Rafael Frongillo, University of Colorado Boulder

Brief Biography: Jessie Finocchiaro (she/her) is a NSF Mathematical Sciences Postdoctoral Research Fellow and Fellow in the Center for Research on Computation and Society (CRCS) at Harvard University, hosted by Yiling Chen. Previously, she completed her PhD in Computer Science at the University of Colorado Boulder under the supervision of Rafael Frongillo and BS in Computer Science and Mathematics at Florida Southern College.

Research Summary: Algorithmic predictions supplement human and algorithmic decision-making in a variety of domains: predicted risk of disease leads to flagged scans for additional review from medical professionals, predicted hyperparameters change physics models for simulator design in engineering, and predicted market trends change investment strategies in quantitative finance. In practice, such algorithmic predictions are often made by training a model which minimizes some *loss* function measuring error, and predictions are used to guide some downstream *decision* or recommendation. Algorithmic predictions experiencing low error can often lead to poor decision-making and recommendations if these algorithms are not able to make “smarter” errors when necessary.

Subtle challenges emerge when incorporating such structure into loss functions. Loss functions often need to balance several desiderata; aligning the loss with decision task is often at odds with designing losses that are computationally tractable to minimize (e.g., convex). Historically, convex losses have been constructed in an ad-hoc manner, and often do not align with the intended decision task. Conversely, feasibility or equity concerns conceptually require algorithm designers to modify the optimized loss, but little work has characterized *how* various constraint formulations change decision-making. *My research agenda examines the bidirectional relationship between algorithm design and decision-making in machine learning and algorithmic economics.*

- (1) *Decision-making to algorithm design.* Given a decision problem, design a “good” loss (part of the machine learning algorithm) that is statistically consistent for the decision problem [1, 3].
- (2) *Algorithm design to decision-making.* Given a fixed algorithm, examine how values embedded into the algorithm change decision-making [2].

Representative Papers:

- [1] An Embedding Framework for the Design and Analysis of Consistent Polyhedral Surrogates (NeurIPS 2019)
with R. Frongillo and B. Waggoner
- [2] Using property elicitation to understand the impacts of fairness constraints (Working paper)
- [3] The Structured Abstain Problem and the Lovász Hinge (COLT 2022)
with R. Frongillo and E. Nueve IV

BAILEY FLANIGAN ([Homepage](#), [CV](#))

Thesis: Supporting Better Direct Democracy via Social Choice, Social Science, and Algorithms ('24)

Advisor: Ariel Procaccia, Harvard School of Engineering and Applied Sciences

Brief Biography: I am a 5th-year PhD student in Computer Science at Carnegie Mellon University. My research combines theoretical and applied techniques from Econ-CS, social choice, and political science to design and support democratic processes that *directly* involve the public in policymaking, usually through *deliberation*. My research has been recognized by an honorable mention for the *INFORMS Doing Good with Good OR* prize, a *SIGCSE Best Paper Award*, and a *Siebel Scholarship*. Outside of research, at CMU I led the creation of *CS-JEDI*, a PhD course on diversity, equity, and inclusion. For this I received CMU-level and department-level service awards. My PhD was funded by a *Hertz Fellowship* and an *NSF GRFP*.

Research Summary: The main goal of my research is to design and support democratic processes that produce greater social benefit, work against patterns of marginalization, and cultivate public trust. I focus especially on processes that *directly* involve the public in policymaking, usually via *democratic deliberation* — a reasoning-based discussion between constituents about a political decision. Deliberative processes are a high-impact and rapidly growing application domain, with deliberative processes like *citizens' assemblies* and *deliberative town halls* now being used to inform political decisions around the world, even at the national level.

Despite the accelerating uptake of deliberative processes, many pressing questions about their principled implementation and impact remain unanswered. My research targets questions in this space, such as:

- Q1* How should participants of deliberative processes be *selected*?
- Q2* What does it mean for a deliberative process to be *representative*?
- Q3* Why — and when — does deliberation *improve democratic outcomes*?

My work so far on *Q1* uses tools from algorithms, convex and integer optimization, and game theory to build randomized selection algorithms that guarantee notions of fairness, transparency, representation, and robustness to manipulation. These algorithms have been published in *Nature* and deployed in real-world deliberative processes by major groups of practitioners. My lines of work on *Q2* and *Q3* are rooted in innovations on standard voting models, and both rely on ideas from social choice and social science. I am now working with political scientists and practitioners to deploy my work on *Q2* within UC Riverside's deliberation platform *Prytaneum*. Likewise, my work on *Q3* is being deployed in ongoing deliberative processes, including a town hall that will guide the reform of Chile's constitution.

Representative Papers:

- [1] Fair Algorithms for Selecting Citizens' Assemblies (*Nature* '21)
with P. Gözl, A. Gupta, B. Hennig, and A.D. Procaccia
- [2] Toward Accounting for Stakes in Voting (*under submission*)
with S. Wang and A.D. Procaccia
- [3] Distortion Under Public-Spirited Voting (*EC* '23)
with S. Wang and A.D. Procaccia

YURI RESENDE FONSECA ([Homepage](#), [CV](#))

Thesis: Learning from Optimal Actions: Theory and Empirical Analysis in Digital Platforms ('24)

Advisor: Omar Besbes, Columbia University, Ilan Lobel, New York University

Brief Biography: I am a Ph.D. candidate in Decision, Risk, and Operations at Columbia University, where I am fortunate to be advised by Omar Besbes and Ilan Lobel from NYU-Stern. I previously obtained an M.Sc. and a D.Sc. in Materials Science from the Military Institute of Engineering in Brazil under the supervision of Carlos Nelson Elias.

Research Summary: My research uses tools from operations research, learning theory, and causal inference to analyze observed decisions (emerging from algorithms or agents) in complex environments and learn from their behavior. I am interested in questions pertaining to structural estimation, mechanism design, and algorithm design.

The first part of my research agenda aims at leveraging structural models and machine learning methods to understand the complex interactions emerging in digital platforms and the implications for their design. In “*Signaling Competition in Two-Sided Markets*,” I collaborated with a large service marketplace to elucidate how and how much workers anticipate competition when applying for short-term jobs. We develop and estimate a structural model that allows workers to incorporate expectations of competition levels and react to information released by the platform. Our structural model is anchored around an equilibrium concept to summarize workers’ interactions at the platform level. I show that by revealing information about competition levels to workers, congestion is reduced, and the expected number of deals in the platform increases.

My second stream of work focuses on single-agent inverse optimization. In particular, I am interested in designing machine learning and optimization methods for learning structural models from revealed preferences. In “*Contextual Inverse Optimization: Offline and Online learning*”, I aim to answer the following fundamental question: What and how fast can one learn based on past actions taken by agents? For the offline setting, we characterize an instance-dependent minimax regret and the corresponding optimal policy. The notion of regret is taken for adversarial new instances and captures the generalization power of the model learned. For the online setting, we improve state-of-the-art results, providing an improvement from $O(\sqrt{T})$ to $O(\log T)$ in the order of the regret. We are able to do so by designing novel online learning algorithms that exploit explicitly the geometry of this class of problems.

Representative Papers:

- [1] Signaling Competition in Two-Sided Markets (EC’23)
with O. Besbes, I. Lobel, and F. Zheng
- [2] Contextual Inverse Optimization: Offline and Online Learning (COLT’21, Operations Research - articles in advance)
with J. Smith, J. Doe, R. Roe, and J. Roe
- [3] “Statistical Learning And Inverse Problems: A Stochastic Gradient Approach (NeurIPS’22) with Y. Saporito

ABHEEK GHOSH ([Homepage](#), [CV](#), [Google Scholar](#))

Thesis: Contests: Equilibrium Analysis, Design, and Learning ('24)

Advisors: Edith Elkind and Paul W. Goldberg, University of Oxford

Brief Biography: I am a fourth-year Ph.D. student at Oxford supported by the Clarendon Fund Scholarship. I spent the summer of 2022 working with Milind Tambe at Google Research. I did my master's at UT-Austin and bachelor's at IIT-Guwahati and was previously advised by Umang Bhaskar (TIFR Mumbai).

Research Summary: My recent research has focused on contest theory. Contests are games where agents compete for valuable rewards by putting in costly and irreversible efforts. Classic examples include the all-pay auction and the Tullock contest. Contests sometimes arise naturally, like competition among firms for drug discovery and patents, among students for college seats, but they may also be explicitly organized, like crowdsourcing, cryptocurrencies, and sports.

The assumptions like the availability of information and full rationality of agents, which are necessary for equilibrium analysis, may not hold in practice for many applications of contests. My ongoing research with Paul Goldberg studies learning dynamics in contests. We show the convergence of best-response dynamics in Tullock contests with homogeneous agents. We also provide almost tight rate of convergence bounds using techniques from convex optimization and randomized algorithms [2] and a novel analysis of a general class of stochastic processes [1]. These papers also show that best-response dynamics may not converge for non-homogeneous agents; for this case, our ongoing work proves convergence for more general dynamics. Another line of research studies the convergence of learning dynamics when the agents only receive bandit feedback, which is the case in many real-life occurrences of Tullock contests. Related future work: I plan to study correlated and course-correlated equilibria and convergence of no-regret dynamics in contests like rank-order allocation (which includes all-pay auction) and Tullock.

My other research in contest theory includes: (i) Designing contests to improve diversity. In one work, we design contests to get higher participation from many agents rather than a very high effort from a few. In another work, our goal is to incentivize a target group to put in more effort. (ii) Equilibrium complexity. We show that computing an equilibrium in a certain class of contests is CLS-complete.

Other topics. In collaboration with a healthcare charity and researchers at Google, we solve a sequential limited resource allocation problem. We give near-optimal algorithms for restless multi-armed bandits using mean-field methods and show improved performance for two healthcare applications [3]. My other projects have focused on social choice theory, in particular, deliberation [4] and voting.

Representative Papers:

- [1] Best-Response Dynamics in Tullock Contests with Convex Costs (WINE 2023)
- [2] Best-Response Dynamics in Lottery Contests (EC 2023) with P. W. Goldberg
- [3] Indexability is Not Enough for Whittle: Improved, Near-Optimal Algorithms for Restless Bandits (AAMAS 2023) with D. Nagaraj, M. Jain, M. Tambe
- [4] Complexity of Deliberative Coalition Formation (AAAI 2022) with E. Elkind, P. W. Goldberg

DENIZALP (DENI) GOKTAS ([Homepage](#), [CV](#), [Google Scholar](#))

Thesis: A Generative Adversarial Theory of Games ('24)

Advisor: Amy Greenwald, Brown University

Brief Biography: Deni is a fifth year Computer Science Ph.D. student at Brown University. His research builds and analyzes multiagent learning algorithms in games and markets with the ultimate goal of building welfare improving technology based on these algorithms. Deni has previously worked as a research scientist intern at JP Morgan's AI research lab as well as Google DeepMind's Game Theory and Multiagent Team, and was a visiting scholar at UC Berkeley's Simons institute.

Research Summary: Over the last two centuries, mathematical economists and game theorists alike, have dedicated a great deal of effort to constructing models of choice that predict the outcome of multiagent interactions. These models can be roughly categorized as general equilibrium models, i.e., market models, and game models. Although both concern preference maximizing agents, the algorithmic literature on general equilibrium and games has evolved mostly independently. I have shown that both classes of models are special cases of two-player zero-sum Stackelberg (i.e., sequential) games, and use this insight to devise efficient algorithms to solve games, and general equilibrium models.

Solving Games: we have recently introduced a class of zero-sum Stackelberg games, along with first-order algorithms to solve such games in polynomial time [1]. We have then shown that any any game is an instance of zero-sum Stackelberg games, and provided polynomial-time first-order methods to approximate equilibria in all Lipschitz-smooth games [2]. In even more recent work, we generalized our results to stochastic games [3].

Solving General Equilibrium Models: we have discovered that a large class of general equilibrium models are equivalent to zero-sum Stackelberg games [1], and have shown an interesting connection between first-order methods to solve these sequential games and age-old auction algorithms [3]. We then generalized this result to stochastic settings, showing that a large class of stochastic general equilibrium models can be seen as zero-sum stochastic Stackelberg games, providing a polynomial-time solution to these games [2].

As our world becomes increasingly more automated, our societies will face an increasing number of multiagent decision making problems. My research aims to solve these problems via the flexible and broad framework of zero-sum Stackelberg games, i.e., generative adversarial optimization. Indeed, the framework and algorithms I introduced have been used to solve problems in automated test generation for autonomous systems, resource allocation, and human-robot interaction [1,2,3].

Representative Papers:

- [1] Convex-Concave Min-Max Stackelberg Games. (NeurIPS'21)
with D. Goktas, and A. Greenwald
- [2] Exploitability Minimization in Games and Beyond. (NeurIPS'22)
with D. Goktas, and A. Greenwald
- [3] Tâtonnement in Homothetic Fisher Markets (EC'23)
with D. Goktas, J. Zhao, and A. Greenwald

YINGKAI LI ([Homepage](#), [CV](#))

Thesis: Approximate Optimality of Simple Mechanisms ('22)

Advisor: Jason Hartline, Northwestern University

Brief Biography: Yingkai is a postdoc at the Cowles Foundation for Research in Economics at Yale University working with Prof. Dirk Bergemann and Prof. Yang Cai. He obtained his PhD from the Department of Computer Science at Northwestern University, where he was advised by Prof. Jason Hartline. He has also interned at Microsoft Research during the summers of 2020 and 2021.

Research Summary: My interests lie broadly at the intersection of computer science and economics, with a focus on mechanism design and information design.

Information elicitation: In many environments, firms or organizations depend on experts to obtain valuable information at a cost to facilitate more informed decision-making. We provide characterizations and compelling economic insights of the optimal mechanisms for incentivizing both the acquisition of costly information and truthful reporting by the agent. We address this in both static [1] and dynamic [2] models. Additionally, we developed polynomial time algorithms for computing these optimal mechanisms in several canonical environments.

Endogenous principal learning: In online marketplaces, platforms can accumulate vast amounts of user data. However, in many cases, these platforms lack the ability to commit to how they will use this collected data in future interactions. This strategic uncertainty can undermine the trust that users place in the platform, which is generally an undesirable outcome for the platforms. In [3], we outline conditions under which the platform consistently favors gathering more user information and conditions under which the platform significantly benefits from abstaining from collecting certain user data.

Simple mechanisms: In many cases, optimal mechanisms tend to be highly complicated and, as a result, impractical for real-world applications. However, in classic auction settings, we show that simple mechanisms can often achieve approximately optimal performance [4], and are robustly optimal when the designer faces uncertainty about the environment [5]. As part of future research, I am interested in the development of simple mechanisms and the evaluation of their performance in environments where the designer can simultaneously design the mechanisms and specify the information to be gathered.

Representative Papers:

- [1] Optimization of Scoring Rules (EC 2022)
with Jason Hartline, Liren Shan, and Yifan Wu
- [2] Optimal Scoring for Dynamic Information Acquisition (working paper)
with Jonathan Libgober
- [3] Mechanism Design with Endogenous Principal Learning (working paper)
with Daniel Clark
- [4] Simple Mechanisms for Non-linear Agents (SODA 2023)
with Yiding Feng and Jason Hartline
- [5] Benchmark Design and Prior-independent Optimization (FOCS 2020)
with Jason Hartline and Aleck Johnsen

ZUN LI ([Homepage](#), [CV](#))

Thesis: Artificial Intelligence Methods for Economic and Computer Games ('23)

Advisor: Michael P. Wellman, University of Michigan, Ann Arbor

Brief Biography: Zun Li is a Ph.D. candidate at University of Michigan, Ann Arbor, advised by Prof. Michael P. Wellman. His research lies on the interface between artificial intelligence and computational economics. He had worked as a research scientist intern at DeepMind Alberta in 2022, working with Dr. Marc Lanctot on game-tree search in general-sum imperfect information games. He also had worked as a software engineer intern at Google display ad auction team.

Research Summary: My Ph.D. research followed precisely the chronological order that most game theory textbooks are organized: the most basic normal-form games are first studied, then are games with incomplete information, and then are dynamical games with imperfect information. The only difference here, though, is that my approaches were more from a computational perspective using practical AI methods, instead of deriving the exact mathematical solutions.

My first work [1] adopted a model-based learning approach to solve normal-form games with many players. By using supervised or unsupervised learning techniques, we can learn a succinct representation (such as clusters or a graph) of the true game using payoff data under some structural hypothesis, instead of storing an $O(NM^N)$ tensor. The computation within the learned game can be much more efficient, and the solutions were experimentally shown well in the true games. In my second work [2], I formulated the equilibrium computation in Bayesian games in a similar way as in Deep RL, where each pure strategy is represented as a neural net, and the utilities come in the form of black-box simulation data. Using natural evolution strategies, I proposed algorithms to compute pure equilibria and mixed equilibria. The methods exploit the symmetry structure of the game and scale well in high-dimensional games. We found that deep neural nets can recover classical analytical solutions in simple games like first- and second-price auctions.

My latest work [3] extended AlphaZero-styled search method to general-sum imperfect information games by replacing MCTS with information-set MCTS, and learning a deep belief network to represent belief states at the root of the search tree. Furthermore, we combine this new search method with policy space response oracle and construct a decision-time AI bot that can conduct test-time search and online Bayesian opponent modeling. We evaluate this bot against humans in a class of negotiation games and found our bot gave comparable social welfare with humans.

Representative Papers:

- [1] Structure Learning for Approximate Solutions of Many-Player Games (AAAI'20)
with M. Wellman
- [2] Evolution Strategies for Approximate Solution of Bayesian Games (AAAI'21)
with M. Wellman
- [3] Search-Improved Game-Theoretic Multiagent Reinforcement Learning in General and Negotiation Games (Extended Abstract) (AAMAS'23)
with M. Lanctot, K. McKee, L. Marris, I. Gemp, D. Hennes, P. Muller, K. Larson, Y. Bachrach and M. Wellman

FRANCISCO MARMOLEJO-COSSÍO ([Homepage](#), [CV](#))

Thesis: Equilibrium Computation in Games and Strategic Aspects of Bitcoin Mining ('20)

Advisor: Paul Goldberg, University of Oxford

Brief Biography: I am a Lecturer and Postdoctoral Fellow at Harvard University hosted by David Parkes. I am also a Research Fellow at Input Output Global (IOG), as well as a Co-organizer of the Mechanism Design for Social Good (MD4SG) research initiative. Previously, I was a Career Development Fellow and Senior Tutor in Computer Science at Balliol College at the University of Oxford, from where I also received a D.Phil. in Computer Science under the supervision of Paul Goldberg, and an M.Sc. in Mathematics and Foundations of Computer Science. I also hold a B.A. in Mathematics (minor Neuroscience) from Harvard University. My current work was recognized as an Exemplary Applied Modeling Track Paper at EC '23, and previous work was nominated for Best Paper at WINE '20.

Research Summary: I work on societally facing challenges, including problems in public health and access to financial services, with the goal of improving access to resources within underserved communities. Beyond methodological advances, I am committed to fostering collaborations between academics and practitioners, in line with the mission of MD4SG, an interdisciplinary research initiative that brings together practitioners and researchers around the world to work towards the aforementioned goal.

My work specifically spans three key problem areas. The first is the issue of resource allocation in resource-constrained communities. I have built, deployed, and piloted algorithmic tools to help decision-makers in Mexican universities allocate scarce testing supplies during the COVID-19 pandemic [1]. In this work, I leveraged techniques from optimization and algorithm design alongside novel modeling paradigms I built through sustained conversations with students, faculty, administration, and testing personnel. Beyond its utility in immediate pandemic response, this work also has potential applications in other general resource-allocation settings. My second specialization area has to do with the disruptive potential of blockchain technology for underserved communities. Given the nascent nature of the technology, my current work has focused on understanding the broader blockchain ecosystem as a precursor to identifying use-cases and adoption barriers. This has involved modeling and studying strategic behavior in decentralized consensus protocols and decentralized finance (DeFi) [2]. In this work, I have partnered with industry partners who deploy and maintain products for a large variety of stakeholders. The third problem area is more theoretical and focuses on informational resources required to compute equilibrium concepts in game theory [3].

Representative Papers:

- [1] Welfare-Maximizing Pooled Testing (EC'23, Exemplary Paper Award)
with S. Finster, M. González-Amador, E. Lock, E. Micha, and A. Procaccia
- [2] Strategic Liquidity Provision in Uniswap v3 (AFT '23)
with Z. Fan, D. Moroz, M. Neuder, R. Rao and D. Parkes
- [3] Learning Convex Partitions and Computing Game-theoretic Equilibria from Best Response Queries (ACM TEAC and WINE '18) with P. Goldberg

VISHNU V. NARAYAN ([Homepage](#), [CV](#))

Thesis: Multi-Item Auctions and Fair Division ('22)

Advisor: Adrian Vetta, McGill University

Brief Biography: Vishnu V. Narayan is a postdoctoral fellow at Tel Aviv University hosted by Michal Feldman. He received his Ph.D. from McGill University in 2022, where he was advised by Adrian Vetta and studied the structure of equilibria in sequential auctions. At present, his main research focus is on fair division; specifically, he is interested in exploring the power of payments in achieving fairness, and in extending fair division results beyond additive valuations and beyond the division of goods. He is also more broadly interested in other research areas within CS theory, including discrete optimization, graph theory and online algorithms, and has published work on a variety of topics. During his Ph.D., he spent a semester as a visiting fellow at Harvard (hosted by Ariel Procaccia). He has a Best Paper award from SAGT 2019 and a Teaching Assistant Award from McGill University.

Research Summary: I am drawn to problems that tie my previous expertise in combinatorics and optimization together with the Econ-CS domain. The main theme of my research is fair division, which asks how to divide a collection of items amongst agents with different preferences in a manner that everyone agrees is fair. The current decade witnessed an explosion of research in this area, but the majority of this activity is in the division of goods amongst agents with additive valuations. One focus of my research is to push past the frontier of additive valuations, and we have recently made significant progress in this area, both for goods [1,2] and chores [3]. Additionally, I am interested in exploring the effects of transfer payments in fair division. In our EC'20 paper [4], we studied the problem of achieving envy-freeness in indivisible-item instances through the use of a subsidy, and gave a tight upper bound on the amount of subsidy sufficient to always eliminate envy, resolving two conjectures in the process. My ongoing research studies the effects of payments in other settings (such as in chore division and scheduling).

One long-term goal is to complete a comprehensive analysis of the power of transfers in fair division. I also hope to make progress on the remaining big open problems in the area (such as the EFX problem); and intend to apply my expertise in fair division beyond item-allocation to other settings (such as two-sided markets, voting, budgeting, and fair algorithms in machine learning).

Representative Papers:

- [1] Fair Division via Quantile Shares [Working Paper]
with Y. Babichenko, M. Feldman, and R. Holzman
- [2] Breaking the Envy Cycle: Best-of-Both-Worlds Guarantees for Subadditive Valuations [Working Paper]
with M. Feldman, S. Murras, and T. Ponitka
- [3] Fair Chore Division under Binary Supermodular Costs [AAMAS'23 ext.abs.]
with S. Barman and P. Verma
- [4] One Dollar Each Eliminates Envy [EC'20]
with J. Brustle, J. Dippel, M. Suzuki, and A. Vetta

ORESTIS PAPADIGENOPOULOS ([Homepage](#), [CV](#), [Google Scholar](#))

Thesis: Online Decision-Making and Learning under Structured Non-Stationarity ('22)

Advisor: Constantine Caramanis, The University of Texas at Austin

Brief Biography: Orestis is a Postdoctoral Research Scientist at the Data Science Institute of Columbia University, hosted by Vineet Goyal and Assaf Zeevi. His research interests fall into the broad area of sequential decision-making under uncertainty, with applications in revenue management and operations. Before joining Columbia, Orestis completed his PhD in Computer Science at The University of Texas at Austin, under the supervision of Constantine Caramanis.

Research Summary: Through my research, my objective is to design efficient and practical algorithms that produce interpretable and (nearly) optimal solutions to sequential decision-making problems. I am particularly interested in providing the necessary framework for developing these algorithms in a systematic way, together with the analytical tools for proving theoretical guarantees on their performance.

My thesis work focuses on modeling and algorithmically leveraging the effect of human behavior in recommendation systems. In order to capture the interplay between chosen actions and altering users' preferences (e.g., satiation or deprivation), I studied non-stationary generalizations of the multi-armed bandit framework, where the reward distribution of each action is a function of the history (see, for example, my work on "recharging" bandits [1]).

In a recent direction, I turned my focus on assortment optimization, namely, the problem of deciding on an collection of goods (e.g., products, services, etc.) to offer to a customer in order to maximize the expected revenue of the seller. In joint work with Goyal, Humair, and Zeevi [2], we studied an online variation of the problem where the potential goods are observed sequentially and assortment decisions are made instantaneously and irrevocably. Assuming prior distributional knowledge on the features of each good, we developed (nearly) optimal threshold-based online policies for standard demand models and feasibility constraints.

In addition to the above, I have worked on various topics of (sequential) decision-making, including prophet inequalities under limited information [3], recurrent optimal-stopping, online learning under complex feedback structures, and resource allocation [4].

Representative Papers:

- [1] Non-Stationary Bandits under Recharging Payoffs: Improved Planning with Sublinear Regret (NeurIPS '22) with C. Caramanis and S. Shakkottai
- [2] MNL-Prophet: Sequential Assortment Selection under Uncertainty (WINE '23) with V. Goyal, S. Humair, and A. Zeevi
- [3] Single-Sample Prophet Inequalities via Greedy-Ordered Selection (SODA '22) with C. Caramanis, P. Dütting, M. Faw, F. Fusco, P. Lazos, S. Leonardi, E. Pountourakis, and R. Reiffenhäuser
- [4] A Constant-Factor Approximation for Generalized Malleable Scheduling under M^h -Concave Processing Speeds (IPCO '22), with D. Fotakis and J. Matuschke

JUSTIN PAYAN ([Homepage](#), [CV](#), [Google Scholar](#))

Thesis: Data-Driven Optimization for Social Good ('24)

Advisor: Yair Zick, University of Massachusetts Amherst

Brief Biography: I am a sixth-year PhD candidate at UMass Amherst, where I work on resource allocation, data science, and peer review systems. During my PhD, I have enjoyed interning at Amazon Alexa as a research intern in 2020 and 2021, Microsoft as a data science intern in 2022, and Adobe Research as a research intern in 2023.

Research Summary: I apply techniques from combinatorial optimization, market design, and data science to high-stakes social problems. My main application is peer review, but I have applied my skills to other problems.

My work ensures that market outcomes are socially desirable, especially when the connections between resource distribution and outcomes of interest are not fully observed. I incorporate techniques from natural language processing, optimization, and machine learning to estimate and optimize over predicted outcomes. Our framework Robust Reviewer Assignment (RRA) assigns reviewers to conference papers by building an uncertainty-aware, predictive model of review quality and ensuring high quality reviews are likely [1]. We are also building algorithms for repeated constrained matching problems, minimizing regret for resource allocation problems over time. A major motivating problem setting is reviewer assignment in ACL Rolling Review.

Because I focus on problems of social importance, I am very interested in translating my algorithms into practice. Our algorithm FairSequence is deployed in the major conference management platform OpenReview. FairSequence assigns reviewers to papers fairly, efficiently, and quickly while remaining flexible to the constantly evolving constraint set required by conference organizers [2]. I am working with the Chief Administrator of IJCAI to build NLP and ML models of review quality, and we are in talks with AAMAS to create a database of reviewer “badges” that will track and reward the best reviewers.

I also apply combinatorial optimization tools to structured NLP and computer vision problems. Transformer models have revolutionized tasks that require natural perception, and can serve as an intuitive user interface. However, many domains have hard structural requirements that cannot be satisfied easily by transformers. At Adobe, we merged computer vision and combinatorial optimization for document structure recognition (ongoing), and at Microsoft we built NLP models for code generation in Excel [3].

Representative Papers:

- [1] Into the Unknown: Assigning Reviewers to Papers with Unknown Affinities (SAGT 2023) with C. Cousins and Y. Zick
- [2] I Will Have Order! Optimizing Orders for Fair Reviewer Assignment (IJCAI 2022) with Y. Zick
- [3] InstructExcel: A Benchmark for Natural Language Instruction in Excel (Findings of EMNLP 2023) with S. Mishra, M. Singh, C. Negreanu, C. Poelitz, C. Baral, S. Roy, R. Chakravarthy, B. van Durme, and E. Nouri

FARZAD POURBABAEE ([Homepage](#), [CV](#))

Thesis: Essays in Venture Capital, Reputation and Learning ('21)

Advisor: Robert M. Anderson and Federico Echenique, UC Berkeley

Brief Biography: Farzad is a postdoctoral fellow at Caltech HSS, working with Federico Echenique and Omer Tamuz. He received his Ph.D. in Economics and M.A. in Statistics from UC Berkeley, advised by Robert M. Anderson.

Research Summary: My research primarily focuses on important economic topics, including mechanism design, experimentation, information economics and social learning. I leverage mathematical tools such as stochastic analysis, high-dimensional statistics, and probability theory to address these subjects.

Mechanism Design: In [1] we study the public-good provision with privacy-aware agents. We propose adding calibrated noise to the individuals' preference messages sent to the planner, thereby protecting their privacy. Using Isoperimetric inequalities on the Boolean hypercube, we analyze the implementability, revenue, social surplus and noise robustness of the optimal provision rules.

Information Economics and Social Learning: In a *misspecified* social learning setting [2], where agents sequentially make decisions by observing a private signal as well as the actions of their predecessors, we show learning outcomes improve if and only if agent' misspecification is *moderate*. In [4], I study a dynamic model of learning and random meetings between a long-lived agent with unknown ability and heterogeneous projects. By applying tools from optimal stopping theory, I find the optimal project selection policy of the agent as a function of her reputation.

Experimentation: My other line of work is related to bandits and strategic experimentation. Specifically, in [3] we investigate the strategic decision between exploring a risky project and exploiting a safe option in a network setting, where individuals observe the past experimentation outcomes of their neighbors in the graph of social connections. In [5], I study the experimentation dynamics of a decision maker in a two-armed bandit setup, where she faces *Knightian* uncertainty regarding the payoff distribution of one arm and thus entertains *Multiplier preferences*. I frame the decision making environment as a two-player differential game against nature in continuous time, and find the optimal experimentation strategy.

Representative Papers:

- [1] Binary Mechanisms under Privacy-Preserving Noise (WINE '23) with Federico Echenique
- [2] The Hazards and Benefits of Condescension in Social Learning (EC '23; Revise and Resubmit at *Theoretical Economics*) with Itai Arieli, Yakov Babichenko, Stephan Müller and Omer Tamuz
- [3] The Impact of Connectivity on the Production and Diffusion of Knowledge (working paper; Informs ADA '22) with Gustavo Manso
- [4] Reputation, Learning and Project Choice in Frictional Economies (Revise and Resubmit at Economic Theory)
- [5] Robust Experimentation in the Continuous Time Bandit Problem (*Economic Theory*, 2020)

ROJIN REZVAN ([Homepage](#), [CV](#), [Google Scholar](#))

Thesis: Simple vs. Optimal in Multi-dimensional Mechanisms ('24)

Advisor: Shuchi Chawla, University of Texas at Austin

Brief Biography: Rojin Rezvan is a fifth-year PhD student at the University of Texas at Austin, advised by Shuchi Chawla. She received her masters degree from the University of Wisconsin-Madison. She is broadly interested in algorithmic game theory and mechanism design. More specifically, she has done research in multi-dimensional mechanism design, fairness in auctions and fair allocation. She is generally interested in the intersection of mechanism design and other fields such as fairness and decentralized systems.

Research Summary: One of the main focuses of my PhD is on the paradigm of "Simple vs. Optimal" in mechanism design for multi-dimensional settings. Multi-item mechanisms can have undesirable properties such as unbounded revenue, *lottery* options in the menu and super-additive pricing function. To circumvent these issues, there are two paths to take: 1) Make some assumptions, such as independence over item value distributions and the buyers' value functions, 2) Examine the validity of the benchmark. The approach we took in [1] and [2] was the latter.

Our proposal is to compare any *simple* mechanism we design to a more realistic benchmark, called "Buy-many". In this setting, it is assumed that each buyer can interact with the menu multiple times. This ensures that super-additive pricing will not happen. The main difference now is while optimal revenue may be unbounded, the gap between revenue of optimal simple mechanisms such as item pricing and optimal buy-many mechanisms is logarithmic in the number of items. In [1], we propose a structure necessary over the item values, with which we will get fine-grained results in terms of approximation and computation of the buy-many revenue via item pricing. In [2], we extend these results and definitions to multi-buyer setting.

I am also interested in algorithmic fairness. In [4], we ask: is it possible that certain allocation algorithms in ad auctions introduce unfairness to the allocations in addition to the data? The answer is yes: an algorithm that always allocates to the highest bidder, such as FPA, could potentially turn minor differences in bids to large differences in allocation. To circumvent the issue, we propose two different algorithms that ensure fairness, while losing a fraction of the optimal social welfare, or consequently revenue. Currently, I am working to extend this work to cases where the advertisers have budgets.

Representative Papers:

- [1] Pricing Ordered Items (STOC 22) with S. Chawla, Y. Teng, C. Tzamos
- [2] Buy-many Mechanisms for Many Unit-demand Buyers (WINE 23) with S. Chawla, Y. Teng, C. Tzamos
- [3] Prophet Secretary Against the Online Optimal (EC 23) with P. Duetting, E. Gergatsouli, Y. Teng, and A. Tsigonias-Dimitriadis
- [4] Individually Fair Auctions for Mutli-Slot Sponsored Search (Best student paper at FORC 22) with C. Chawla, N. Sauerberg

ABHIN SHAH ([Homepage](#), [CV](#), [Google Scholar](#))

Thesis: Data Rich Causal Inference ('24)

Advisors: Devavrat Shah, Gregory W. Wornell, and Alberto Abadie, MIT

Brief Biography: Abhin Shah is a final-year Ph.D. student in the department of Electrical Engineering and Computer Science at the Massachusetts Institute of Technology (MIT), where he works with Prof. Devavrat Shah, Prof. Greg Wornell, and Prof. Alberto Abadie. He is a recipient of MIT's Jacobs Presidential Fellowship.

Research Summary: My research develops methods for causal inference from observational data with the goal of personalized, data-driven decision-making. In particular, I address individual-related and distributional questions in causal inference, e.g., what will be the distribution of a consumer's behavior if we expose them to a product? These levels of analysis offer a more nuanced understanding compared to traditional approaches focused on population-related mean questions.

To tackle confounding in observational data, I leverage tools and frameworks from machine learning and statistics. Below, I provide a summary of my research across three threads, each exploiting a different structure to account for confounding.

- (1) *Structure in distribution:* I develop computationally efficient estimators for learning exponential family distributions and use this framework to model unobserved confounding. My work estimates individual-level means of outcomes with just one sample per individual. Integrating this framework with matrix completion, we infer individual-level distributions of outcomes.
- (2) *Structure in equation model:* I use matrix completion to exploit low-dimensional relationship between outcomes and unobserved factors, as well as interventions and unobserved factors, to provide doubly robust estimates of individual-level means of outcomes. These estimates remain accurate even if either outcomes-factors or interventions-factors relationship is mis-specified.
- (3) *Structure in causal graph:* I design data-driven conditional independence tests to infer population-level distribution of the outcome. These tests identify subsets of the observed data that account for unobserved factors, with limited knowledge of the causal generative graph.

My long-term mission is to construct robust and equitable models. While causality is key to robustness, algorithmic fairness plays a pivotal role in ensuring equity. In pursuit of this mission, my research has eliminated disparities across protected groups in applications such as healthcare and criminal-justice by (i) introducing a new notion of fairness for models that predict selectively and (ii) ensuring fairness when protected attribute data is uncertain.

Representative Papers:

- [1] On Counterfactual Inference with Unobserved Confounding (under review at Operations Research) with R. Dwivedi, D. Shah, and G. W. Wornell
- [2] Doubly Robust Causal Inference with Latent Factor Models (working paper) with A. Abadie, A. Agarwal, and R. Dwivedi
- [3] Front-door Adjustment Beyond Markov Equivalence with Limited Graph Knowledge (NeurIPS 2023) with K. Shanmugam and M. Kocaoglu

KANGNING WANG ([Homepage](#), [CV](#), [Google Scholar](#))

Thesis: Approximations for Economic Efficiency and Fairness ('22)

Advisor: Kamesh Munagala, Duke University

Brief Biography: Kangning is a Motwani Postdoctoral Fellow at Stanford University, hosted by Moses Charikar and Aviad Rubinstein. He earned his Ph.D. in Computer Science from Duke University, advised by Kamesh Munagala. He was a J.P. Morgan Research Fellow of the program *Data-Driven Decision Processes* at the Simons Institute, UC Berkeley. He interned twice at Google Research, hosted by Jieming Mao, Renato Paes Leme, and Aranyak Mehta. His work has been recognized by an ACM SIGecom Doctoral Dissertation Award Honorable Mention, a Duke CS Best Dissertation Award, and the WINE 2018 Best Paper Award.

Research Summary: I work broadly in Economics and Computation, with a focus on mechanism design, social choice, information design, and algorithmic fairness. I am particularly interested in developing economic solutions with approximation guarantees for economic objectives such as utility, revenue, efficiency, and fairness.

Bilateral trade is a common economic scenario with a rich literature. The celebrated Myerson-Satterthwaite impossibility theorem shows that bilateral trade generally cannot be efficient under incentives. However, in [1], we reveal that simple mechanisms can always be approximately efficient, answering a prominent open question. Our other works explore the impact of budget constraints, valuation correlation, lack of priors, and impatience on pricing, auctions, and bilateral trade.

In [2], we use a classical price discrimination model to explore how an entity with information about consumer valuations can increase consumer welfare by *persuading* the seller to set personalized prices. We show that, surprisingly, there exists a solution that approximately maximizes all “reasonable” welfare functions simultaneously. Our other works extend this classical price discrimination model to auctions and bilateral trade.

Metric distortion in social choice is a well-studied framework that measures the efficiency of voting rules. In our works including [3], we design new simple voting rules that broke long-standing efficiency barriers (for deterministic rules and for possibly randomized ones), answering frequently asked open questions.

Committee selection and participatory budgeting are common democratic scenarios that require fair solutions. The *core* is often considered to provide the ultimate form of proportionality guarantees. In many of our settings including [4], even though core solutions may not exist, we show that natural relaxations always do.

Representative Papers:

- [1] Approximately Efficient Bilateral Trade (STOC 2022)
with Y. Deng, J. Mao, and B. Sivan
- [2] Fair Price Discrimination (SODA 2024)
with S. Banerjee, K. Munagala, and Y. Shen
- [3] Breaking the Metric Voting Distortion Barrier (SODA 2024)
with M. Charikar, P. Ramakrishnan, and H. Wu
- [4] Approximately Stable Committee Selection (STOC 2020)
with Z. Jiang and K. Munagala

LILY XU ([Homepage](#), [CV](#), [Google Scholar](#))

Thesis: High-stakes decisions from low-quality data: AI decision-making for planetary health ('24)

Advisor: Milind Tambe, Harvard

Brief Biography: I am a CS PhD student at Harvard developing methods in machine learning, sequential planning, and causal inference for planetary health challenges, particularly biodiversity conservation and public health. My research has been recognized with AAAI best paper runner-up, the INFORMS Doing Good with Good OR award, a Google PhD Fellowship, and a Siebel Scholarship. I also co-organize the Mechanism Design for Social Good (MD4SG) research initiative.

Research Summary: I develop and deploy machine learning methods to make reliable decisions in high-stakes settings when data are incomplete. Guided by research questions that emerge from my close collaboration with the public sector, my work enables practitioners to take efficient, robust actions necessary for planetary health. I have worked closely with NGOs and government stakeholders to deploy AI for on-the-ground conservation and maternal healthcare. My research spans:

learning under uncertainty. Online learning enables us to proactively collect more data to improve our models, but may lead to unnecessary exploration. I draw inspiration from immersion in the domains with which I work to integrate problem structure into algorithm design, reducing exploration to achieve higher reward more quickly for multi-armed bandits [1] and reinforcement learning [3].

robust, efficient sequential planning. Resource allocation problems involve multi-step sequential decisions and combinatorial actions, introducing exponentially large action spaces and NP-hard optimization. I develop algorithms with provably strong guarantees to make these challenging problems more tractable, integrating advances from game theory and mixed-integer programming [2].

causal inference for impact evaluation. Causal inference is hard when we cannot conduct RCTs and the available data is messy. I show that machine learning can help overcome these missing data challenges, studying ranger patrol data from a national park in Uganda with one of the highest levels of poaching in the world. Our results provide the first causal evidence for poaching deterrence, showing that ranger patrols reduce poaching by 46%.

Key to my work is my commitment to **deployment**. I partner closely with stakeholders to identify bottlenecks in existing algorithmic solutions and deploy AI to achieve measurable impact on the ground. My work on predicting poaching hotspots has been deployed in multiple countries and is being scaled to 1,200 protected areas worldwide through integration with SMART conservation software.

Representative Papers:

- [1] Dual-Mandate Patrols: Multi-Armed Bandits for Green Security (AAAI 2021) Lily Xu, Elizabeth Bondi, Fei Fang, Andrew Perrault, Kai Wang, Milind Tambe
- [2] Robust Reinforcement Learning Under Minimax Regret for Green Security (UAI 2021) Lily Xu, Andrew Perrault, Fei Fang, Haipeng Chen, Milind Tambe
- [3] Optimistic Whittle Index Policy: Online Learning for Restless Bandits (AAAI 2023) Kai Wang*, Lily Xu*, Aparna Taneja, Milind Tambe

AVIV YAISH ([Homepage](#), [CV](#), [Google Scholar](#))

Thesis: The Security and Economics of Cryptocurrencies ('24)

Advisor: Aviv Zohar, The Hebrew University

Brief Biography: Aviv is a Ph.D. candidate in the Computer Science department at the Hebrew University. His research delves into the intricate relationship between the economics and security of distributed systems. Among other honors, Aviv received the four-year merit-based Ze'ev Jabotinsky Fellowship for Ph.D. students, the Hebrew University's rector award for first-in-class computer science M.Sc. students, and the Austria-Israel Academic Network Innsbruck visiting researcher fellowship. Between the second year of his M.Sc. and until '23, Aviv served as a lecturer for two undergraduate courses, and has received an award for his teaching. Aviv was also a research associate at Matter Labs, and a visiting researcher at the University of Innsbruck, having visited twice during the spring and autumn of '23.

Research Summary: Cryptocurrencies use decentralized mechanisms to facilitate transactions amounting to a daily volume of billions of USD. Consequently, the actors involved in operating cryptocurrencies (i.e., *miners*) may have large monetary incentives to manipulate these mechanisms. Thus, the fundamental question driving my research is: *how can one design mechanisms that align the incentives of cryptocurrency actors, and ensure the secure and smooth operation of the system?*

To answer this question, I adopt a blend of theoretical and practical approaches. The latter, in particular, allow my work to produce a positive real-world impact on the design of popular cryptocurrency mechanisms.

In [1,3], my colleagues and I formalize and analyze novel strategic deviations for miners. In [1], we analyze the “clash” of incentives created when financial applications operate on top of a cryptocurrency: profits extracted from applications can incentivize miners to manipulate cryptocurrency mechanisms. In [3] we unveil two major results: alongside a novel class of miner deviations, we also present the *first* evidence of miners of a large cryptocurrency deviating strategically in the wild. Specifically, we show that Ethereum's mechanism was under attack for two years.

In [2], we show that the transaction fee mechanism used by Ethereum is flawed: it only compensates miners for processing transactions that are included in the blockchain. Thus, adversarial actors can lower transaction throughput by creating transactions that are recognized as ineligible for inclusion only at the end of their execution, without fully compensating miners for their work. We also show that heuristics used to avoid this issue can be circumvented.

In [1,2,3], we analyze modifications to the discussed mechanisms that diminish the profitability of our manipulations. These were disclosed to the relevant cryptocurrency foundations and companies, leading to changes to their mechanisms.

Representative Papers:

- [1] Blockchain Stretching & Squeezing: Manipulating Time for Your Best Interest (EC'22) with S. Tochner, and A. Zohar
- [2] Speculative Denial-of-Service Attacks in Ethereum with K. Qin, L. Zhou, A. Zohar, and A. Gervais
- [3] Uncle Maker: (Time)Stamping Out the Competition in Ethereum (CCS'23) with G. Stern, and A. Zohar

Index

- AI for social impact
 - Lily Xu, 32
- algorithm design
 - Yeganeh Alimohammadi, 7
 - Bailey Flanigan, 18
 - Orestis Papadigenopoulos, 26
- algorithmic decision-making
 - Jessie Finocchiaro, 17
- algorithmic fairness
 - Kate Donahue, 16
 - Abhin Shah, 30
 - Kangning Wang, 31
- algorithmic game theory
 - Bailey Flanigan, 18
- algorithms and uncertainty
 - Hedyeh Beyhaghi, 8
- applied probability
 - Yeganeh Alimohammadi, 7
- auctions
 - Michael Curry, 12
 - Vishnu V. Narayan, 25
- blockchain
 - Francisco Marmolejo-Cossío, 24
 - Aviv Yaish, 33
- causal inference
 - Abhin Shah, 30
- computational social science
 - Zhun Deng, 15
- consistency
 - Jessie Finocchiaro, 17
- contests
 - Abheek Ghosh, 20
- decentralized finance
 - Aviv Yaish, 33
- deliberation
 - Bailey Flanigan, 18
- differentiable economics
 - Michael Curry, 12
- digital platforms
 - Yuri Resende Fonseca, 19
- discrete optimization
 - Vishnu V. Narayan, 25
- dynamics in games
 - Abheek Ghosh, 20
- equilibrium complexity
 - Abheek Ghosh, 20
- equilibrium computation
 - Denizalp (Deni) Goktas, 21
 - Zun Li, 23
- fair auction design
 - Rojin Rezvan, 29
- fair division
 - Johannes Brüstle, 9
 - Vishnu V. Narayan, 25
 - Justin Payan, 27
- game theory
 - Kate Donahue, 16
 - Farzad Pourbabaee, 28
 - Rojin Rezvan, 29
- general equilibrium theory
 - Denizalp (Deni) Goktas, 21
- hard evidence
 - Sulagna Dasgupta, 13
- information design
 - Yingkai Li, 22
 - Kangning Wang, 31
- information economics
 - Farzad Pourbabaee, 28
- learning theory
 - Hedyeh Beyhaghi, 8
- machine learning
 - Michael Curry, 12
 - Kate Donahue, 16
- market design
 - Justin Payan, 27
- matching
 - Tom Demeulemeester, 14
- mechanism design
 - Hedyeh Beyhaghi, 8
 - Johannes Brüstle, 9
 - Linda Cai, 10

- Michael Curry, 12
- Sulagna Dasgupta, 13
- Abheek Ghosh, 20
- Yingkai Li, 22
- Farzad Pourbabaee, 28
- Rojin Rezvan, 29
- Kangning Wang, 31
- Aviv Yaish, 33
- mechanism design for social good
 - Francisco Marmolejo-Cossío, 24
- ML
 - Lily Xu, 32
- ML optimization
 - Johannes Brüstle, 9
- multi-agent reinforcement learning
 - Zun Li, 23
- multi-armed bandits
 - Lily Xu, 32
- multiagent learning
 - Denizalp (Deni) Goktas, 21
- natural language processing
 - Justin Payan, 27
- network science
 - Yeganeh Alimohammadi, 7
- online ad auctions
 - Yurong Chen, 11
- online algorithms
 - Linda Cai, 10
 - Yingkai Li, 22
 - Vishnu V. Narayan, 25
- online decision-making
 - Johannes Brüstle, 9
 - Orestis Papadigenopoulos, 26
- online learning
 - Yuri Resende Fonseca, 19
 - Orestis Papadigenopoulos, 26
- operations research
 - Yeganeh Alimohammadi, 7
 - Tom Demeulemeester, 14
- personalized decision-making
 - Abhin Shah, 30
- private information games
 - Yurong Chen, 11
- randomization
 - Tom Demeulemeester, 14
- reinforcement learning
 - Lily Xu, 32
- resource allocation
 - Francisco Marmolejo-Cossío, 24
- responsible ML
 - Zhun Deng, 15
- revealed preferences
 - Yuri Resende Fonseca, 19
- revenue management
 - Orestis Papadigenopoulos, 26
- security
 - Aviv Yaish, 33
- social choice
 - Bailey Flanigan, 18
 - Kangning Wang, 31
- social good
 - Justin Payan, 27
- Stackelberg equilibria
 - Yurong Chen, 11
- statistics
 - Zhun Deng, 15
 - Farzad Pourbabaee, 28
- strategic communication
 - Sulagna Dasgupta, 13
- strategic learning
 - Linda Cai, 10
- Surrogate loss functions
 - Jessie Finocchiaro, 17
- test design
 - Sulagna Dasgupta, 13

Market Design Job Market Candidate Profiles 2024

Inspired by the SIGecom Exchanges' annual survey of job market candidates,¹ this is the fourth annual collection of profiles of the junior faculty job market candidates of the market design community. The eleven candidates are listed alphabetically. Along with information regarding the candidate's bio, job market paper, other representative papers, and short research summary, each profile also contains links to the candidate's homepage and CV.

We dedicate this effort in memory of all the innocents who were murdered or otherwise harmed in the terrorist attack on Israel on Oct 7, and all the innocents who were killed or otherwise harmed in its ongoing aftermath, and with the hope that all those abducted—children, women, and men—will be back home by the end of this job market season.

—Yannai A. Gonczarowski, Assaf Romm, and Ran Shorrer

Contents

Bilgin, Günnur Ege	37
Demeulemeester, Tom	37
Hahm, Dong Woo	37
Mass, Helene	38
Monjoie, Leopold	38
Nikzad, Afshin	39
Rubbini, Giacomo	39
Saritaç, Ömer	40
Schlom, Christoph	40
Sweat, Kurt	41
Tamura, Yuki	41

¹Starting with: DUGHMI, S., GKATZELIS, V., AND HARTLINE, J. 2015. SIGecom job market candidate profiles 2016. *ACM SIGecom Exchanges 14*, 1, 2–40.

BILGIN, GÜNNUR EGE ([Homepage](#), [CV](#))

Job market paper: Decentralized Many-to-One Matching with Random Search

Advisors: Stephan Lauermaun and Daniel Krähmer

Other fields: Matching theory, Political economy

PhD: University of Bonn, Economics (Expected: 2024)

Short research summary: I analyze finite many-to-one matching markets within a decentralized search and matching framework. As time gets costless, I show there might be no Markovian strategy profile that guarantees stable matchings, and unstable matchings might be guaranteed under equilibrium.

Other papers:

[1] Placement with Assignment Guarantees and Semi-Flexible Capacities, With Aygün, Orhan.

[2] Voting under Salience Bias and Strategic Extremism. With Destan, Cavit Gökem.

DEMEULEMEESTER, TOM ([Homepage](#), [CV](#))

Job market paper: Fair Integer Programming under Dichotomous Preferences. With Goossens, Dries, Hermans, Ben, and Leus, Roel.

Advisors: Roel Leus and Dries Goossens

Other fields: Operations research, Algorithmic game theory, Randomization

PhD: KU Leuven, Research Centre for Operations Research and Statistics (Expected: 2024)

Short research summary: I am interested in problems on the intersection of economics and computation, and have worked on topics such as matching, coalition formation and voting. I have a particular interest in settings where fairness can only be obtained by introducing randomization. In my job market paper, for example, I study how one can fairly return one of the (possibly many) optimal solutions of an integer programming formulation. As illustrated on the kidney exchange problem, the proposed algorithms significantly increase the individual fairness of the affected agents when integer programs are used to make high-impact decisions.

Other papers:

[1] A Pessimist's Approach to One-Sided Matching, *European Journal of Operational Research* (2023). With Goossens, Dries, Hermans, Ben and Leus, Roel.

[2] Rawlsian Assignments. With Pereyra, Juan S.

HAHM, DONG WOO ([Homepage](#), [CV](#))

Job market paper: Leveraging Uncertainties to Infer Preferences: Robust Analysis of School Choice. With Che, Yeon-Koo and He, YingHua.

Advisors: Yeon-Koo Che and Miguel Urquiola

Other fields: Economics of education, Industrial organization

PhD: Columbia University, Economics (2022)

Post-doc: University of Southern California, Economics (2022–2024)

Short research summary: I am an applied microeconomist with expertise in the fields of Economics of Education, Market Design, and Industrial Organization, with a specific emphasis on empirical school choice. I combine design-based reduced-form analysis with tools from market design and empirical IO to gain insights into students’ decision-making processes throughout their academic journeys.

Other papers:

[1] A Dynamic Framework of School Choice: Effects of Middle Schools on High School Choice, *EC* 2022. With Park, Minseon.

[2] Sorting, Commuting, and School Choice, *EC* 2023. With Park, Minseon.

MASS, HELENE ([Homepage](#), [CV](#))

Job market paper: Optimal Testing in Disclosure Games. With Avi Lichtig.

Advisor: Achim Wambach

Other fields: Microeconomic theory, Information economics, Auction theory

PhD: University of Cologne, Economics (2018)

Post-doc: University of Bonn, Economics (2018–2024)

Short research summary: My research focuses on auction theory with applications to procurement in uncertain environments and on information economics with applications in organizational economics, social learning, and disclosure games. In my JMP I derive the optimal information-gathering process for a regulator who aims at being informed on her own but also aims at incentivizing disclosure from informed agents.

Other papers:

[1] Worst-Case Equilibria in First-Price Auctions, *Theoretical Economics* (2023). With Gretschko, Vitali.

[2] Two-Dimensional Information Acquisition in Social Learning, *Journal of Economic Theory* (2022). With Bobkova, Nina.

MONJOIE, LEOPOLD ([Homepage](#), [CV](#))

Job market paper: Designing Markets for Reliability with Incomplete Information

Advisor: Fabien Roques

Other fields: Environmental and ecological economics, Industrial organisation

PhD: Paris-Dauphine University, Economics (Expected: 2024)

Short research summary: My research sits at the intersection of market design and industrial organization, focusing on energy and environmental markets. My

JMP examines the challenges of allocating a good subject to capacity constraints when considering consumer preferences and investment decisions. The lack of complete information about consumer utility and constraints on the implementable mechanism implies that the optimal allocation can lead to discriminating against consumers based on their types and that discrimination depends on the level of investment considered. It has significant welfare and distributive implications, particularly in the context of the ongoing energy transition that demands substantial investments in clean technologies.

Other papers:

- [1] Securing Investment for Electricity Markets. How to Design the Demand Side of Capacity Markets?
- [2] Designing Efficient Capacity Mechanisms: Bidding Behavior and Product Definition

NIKZAD, AFSHIN ([Homepage](#), [CV](#))

Job market paper: Optimal Allocation via Waitlists: Simplicity through Information Design

Advisors: Al Roth and Itai Ashlagi

Other fields: Mechanism design, Matching markets

PhD: Stanford University, Economics (2018); Stanford University, Operations Research (2018)

Short research summary: I focus on markets where traditional approaches that leverage market-clearing prices cannot achieve social objectives such as social welfare and equality. Levers other than prices become necessary. My research provides guidelines for designing markets that achieve these objectives using the following two levers: the underlying market mechanism which determines in what way the market is cleared, and the information elicited from and provided to the participants. Broadly, I take this approach in the context of (i) improving efficiency in matching markets, and (ii) market designs for fairness and distributional equality.

Other papers:

- [1] Efficiency and Equity in Dynamic Matching: Extreme Waitlist Policies, *Management Science* (2023). With Strack, Philipp.
- [2] Matching in Dynamic Imbalanced Markets, *The Review of Economic Studies* (2022)

RUBBINI, GIACOMO ([Homepage](#), [CV](#))

Job market paper: Mechanism Design without Rational Expectations

Advisors: Roberto Serrano, Geoffroy De Clippel, Kareen Rozen, and Pedro Dal Bó

Other fields: Mechanism design, Behavioral economics, Experimental economics

PhD: Brown University, Economics (Expected: 2024)

Short research summary: I am a microeconomic theorist interested in whether classic results in mechanism design hold even when agents are boundedly rational. In my Job Market Paper, I show that incentive compatibility is still required for full implementation whenever agents can accurately predict the outcome of the implementing mechanism, even if their beliefs about their opponents' strategies are incorrect.

Other paper:

[1] The Revelation Principle without Rational Expectations

SARITAÇ, ÖMER ([Homepage](#), [CV](#))

Job market paper: Centralized versus Decentralized Pricing Control for Dynamic Matching Platforms

Advisor: Ali Aouad

Other fields: Design and analysis of matching algorithms

PhD: London Business School, Management Science and Operations (Expected: 2024)

Short research summary: My research focuses on the design of pricing and matching systems for service platforms within the gig economy. Through my research, I aim to: (i) design computationally efficient and practical market algorithms, and (ii) improve social welfare and equity of market outcomes. I value maintaining strong connections with the industry and adopting an application-oriented research style.

Other paper:

[1] Dynamic Stochastic Matching Under Limited Time, *Operations Research* (2022). With Aouad, Ali.

SCHLOM, CHRISTOPH ([Homepage](#), [CV](#))

Job market paper: Price Distribution Regulation

Advisor: Phil Reny

Other fields: Regulation, Mechanism design

PhD: Chicago Booth, Economics (Expected: 2024)

Short research summary: My job market paper studies the optimal price-based regulation of monopolists who use product quality to discriminate between consumers (Mussa-Rosen monopolists). The price distribution regulation that I study highlights an important friction when regulation is “one-sided” (i.e., only touches price or quality): namely, that a regulator cannot direct surplus to low willingness-to-pay consumer types. One surprising implication of this fact is that, when the regulator separates surplus for different consumer types, if the monopolist's participation constraint does not bind under the optimal regulation, a more equity-focused regulator will pursue a less severe regulation, which harms all consumer types and helps the monopolist.

Other papers:

- [1] Regulating Platform Procurement and Self-Production
- [2] Sharpening Winkler’s Extreme Point Theorem, and Economic Applications

SWEAT, KURT ([Homepage](#), [CV](#))

Job market paper: Endogenous Priority in Centralized Matching Markets: The Design of the Heart Transplant Waitlist

Advisors: Alvin Roth, Frank Wolak, Paulo Somaini, Itai Ashlagi, and Han Hong

Other fields: Health economics, Industrial organization

PhD: Stanford University, Economics (Expected: 2024)

Short research summary: I am interested in studying market design with applications in healthcare using econometric models grounded in economic theory. My job market paper studies the usage of treatments for end-stage heart failure to assign priority in the heart transplant waitlist. I estimate doctors’ preferences over treatments/transplants using administrative data. I show that patients who receive different treatments in response to priority have worse health outcomes mainly because this increases the option value of declining offers, so that these patients spend more time waiting and are more likely to die on the waitlist.

TAMURA, YUKI ([Homepage](#), [CV](#))

Job market paper: Obviously Strategy-proof Rules for Object Reallocation

Advisor: William Thomson

Other fields: Political economy

PhD: University of Rochester, Economics (2021)

Post-doc: New York University, Abu Dhabi (2021–2024)

Short research summary: I am interested in market design with a focus on resource allocation problems. My research agenda is to express social objectives formally, to understand their implications as completely as possible, and to design procedures that meet the objectives. In my job market paper, I study problems where resources are indivisible. I develop procedures that satisfy efficiency and fairness, and that are immune to strategic behavior in a strong sense.

Other papers:

- [1] The Crawler: Three Equivalence Results for Object (Re)allocation Problems when Preferences are Single-Peaked, *Journal of Economic Theory* (2022). With Hosseini, Hadi.
- [2] Object Reallocation Problems with Single-Dipped Preferences, *Games and Economic Behavior* (2023)

SIGecom Winter Meeting 2024 Highlights

AGHAHEYBAT MAMMADOV

Penn State University

and

EMILY RYU

Cornell University

and

ROBERTO SAITTO

Stanford University

***Aghaheybat Mammadov** is a PhD candidate in Economics at Penn State University, advised by Nima Haghpanah and Hadi Hosseini. Within the Economics Department and the AI Lab (College of IST), he studies problems at the intersection of Economics and AI, focusing on resource allocation and matching markets through analytic, algorithmic, and quantitative perspectives. He holds a BSc in Economics from ADA University and an MA in Economics from Bilkent University.*

***Emily Ryu** is a third year PhD student in Computer Science at Cornell University, advised by Éva Tardos and Jon Kleinberg. Her research interests span algorithmic game theory, combinatorial optimization, and networks, particularly with more realistic models of behavioral and cognitive constraints. Before Cornell, she graduated from Princeton University with a B.A. in Chemistry and minors in applied math and computer science.*

***Roberto Saitto** is a fourth-year PhD student in Economics at Stanford University, advised by Paul Milgrom and Matt Jackson. His research interests are economic theory and market design, with a particular focus on simplicity in mechanism design.*

The fourth annual ACM SIGecom Winter Meeting took place on February 15, 2024. Organized by Sigal Oren and Ran Shorrer, this year's meeting brought together researchers from economics, computer science, and adjacent fields to focus behavioral models. The virtual meeting took the form of a workshop including talks and presentations from leading experts on various directions, as well as a fireside chat on getting into the research space and exciting opportunities that lie ahead. We share some highlights from the 2024 Winter Meeting, and additional insights from follow-up interviews with the speakers.

1. INTRODUCTORY TALKS

1.1 Jon Kleinberg - Behavioral Agents in Algorithmic Environments

Consider the everyday scenarios where algorithms play a pivotal role: from recommending movies or social media content based on a user's digital footprint, to aiding hiring committees in job or school applications. In both the online and offline realms, individuals are mapped into vectors of features, which are then processed by decision-makers—algorithms, humans, or both—to predict outcomes. Moreover, algorithms can step beyond mere tools to become partners. This dynamic is evident in generative AI assistants, semi-autonomous vehicles, and even medical diagnostic tools. **Jon Kleinberg** (Cornell University) posed a crucial question in the first introductory talk of the 2024 SIGecom Winter Meeting: To what extent do algorithms need to incorporate behavioral models to truly enhance the welfare of users across these diverse environments? How can we ensure that these intelligent systems do more good than harm? Jon's exploration into the intersection of computational economics and human behavior makes us to rethink the role of algorithms in shaping our decisions and futures.

One model of agents with behavioral biases addresses how agents make decisions over time with a *present bias*, such as procrastination or task abandonment [Kleinberg et al. 2016]. Agents navigate a graph from a starting node s to a terminal node t , aiming to maximize rewards while minimizing costs. However, due to present bias, costs of immediate actions are perceived as higher, scaled by a factor $b > 1$. This bias often leads to agents abandoning their intended paths. To mitigate this, algorithms can restrict certain choices, effectively narrowing down the decision paths and helping agents commit to more optimal routes. This approach illustrates how strategic limitations in choice architecture can assist agents in overcoming inherent biases for better decision-making outcomes.

Another model examines how behavioral biases affect decisions in the consumption of social media feeds [Kleinberg et al. 2023]. A user consumes a content feed of items, and the main decision is whether to continue to the next item. A user has two sets of preferences which are referred to as System 1 and System 2. System 1 is impulsive and acts first, while System 2 uses long-range planning after System 1 chooses not to act. This model questions the reliability of session length as a metric for user welfare. Longer engagement might suggest higher utility if System 1 dominates, but this is not always the true story if System 2 is taken into account. Variability in user's impulsivity can distort this metric, as long sessions could be driven by either genuine interest or mere impulsiveness. This model underscores the need for careful interpretation of user data and design choices in digital platforms to truly enhance user experience and welfare.

The rest of the talk was on an environment where algorithms operate as *partners*. Jon discussed the design of algorithms that remain within the comprehension limits of human partners [McIlroy-Young et al. 2020]. He highlighted "Maia", a chess engine trained on human games, designed to mimic human moves at specific player rankings, known as Elo ratings, offering a seamless transition if a human needs to take over. This contrasts with more sophisticated engines like Leela and Stockfish, which can confound human partners with complex moves. In a unique chess experiment featuring teams composed of both humans and engines with no inter-team

communication, teams often faltered when humans couldn't understand or follow the engines' advanced strategies. However, a modified version of Leela named "PartnerBot" was designed to select moves that were intelligible to both human players and Maia, enhancing team performance. This experiment underscores the importance of transparency and intelligibility in collaborative environments where algorithms and humans interact, ensuring smoother hand-offs and more effective teamwork.

1.2 Ori Heffetz - EBRD and Deferred Acceptance

In his intriguing talk, **Ori Heffetz** (Cornell University) shed light on a puzzling phenomenon: Why do individuals misreport their preferences even when participating in strategy-proof mechanisms such as Deferred Acceptance (DA)? Despite the theoretical assurance that participants in DA—such as students in school admissions and residents in hospital-resident matching—have no incentive to manipulate their preferences, empirical evidence suggests otherwise. Participants appear to make choices that are counter-intuitive or dominated. Ori proposed that incorporating *expectations-based reference-dependent (EBRD) preferences* might be the key to understanding and rectifying this inconsistency in behavior [Dreyfuss et al. 2022].

The literature offers various explanations for this phenomenon, ranging from a lack of understanding that DA is indeed strategy-proof [Li 2017] to misconceptions about the mechanism itself [Gonczarowski et al. 2023]. Another proposal is that altruism or a preference for higher-ranked choices might drive these behaviors [Meisner 2023]. Ori argued that applicants might be influenced not just by the schools they attend but also by expectations-based reference points, known as EBRD preferences. This model incorporates not only traditional utility from consumption but also the utility derived from "news"—the variance between new and old beliefs about consumption. The framework effectively explains why applicants might distort or omit their true preferences, often due to the fear of potential disappointment.

This theory is supported by Ori's incentivized experiment, in which he analyzed four configurations of matching environments. These configurations varied depending on whether students were on the proposing or receiving side, and whether they decided which schools to apply to statically at the beginning of the mechanism or dynamically during its operation. Ori showed that in all of these matching variants, except when students receive offers from schools dynamically (where the "news" effect is absent), the EBRD model comprehensively explained the non-intuitive behaviors or misrepresentations by applicants, starkly contrasting with traditional preference theories. This insight challenges us to rethink how we design and interpret the results of matching mechanisms, highlighting the complex interplay between expectations and decision-making.

2. FIRESIDE CHAT WITH NOAM NISAN AND ALVIN ROTH

The 2024 Winter Meeting featured a dynamic Q&A session, the Fireside Chat. In this event, two renowned professors, **Noam Nisan** (Hebrew University), a distinguished computer scientist specializing in computational complexity theory and algorithmic game theory, and **Alvin Roth** (Stanford University), a prominent economist and a Nobel Prize winner especially known for his work in market design, engaged with an audience of graduate students and established scholars. Noam and Alvin offered valuable insights into their PhD journeys and academic careers. They explored the challenges they faced, the learning opportunities they encountered, and how they have grown within the academic community.

Could you describe what a typical week looked like for you when you were a PhD student?

Noam. I participated in an extremely rigorous PhD program at Stanford University. The experience was highly demanding for each week, providing a deep dive into advanced research methodologies and critical thinking, which significantly honed my analytical and academic skills.

Alvin. I met with my advisor every week. The meeting lasted an hour, with the initial half dedicated to discussing the reasons behind the limited progress on my projects. The latter half focused on my advisor recommending specific papers for review, emphasizing the importance of understanding their content to aid in my research endeavors. This structure balanced addressing challenges while also providing constructive guidance for moving forward academically.

What are your thoughts on the advice that selecting a good advisor is more important than choosing a specific research topic for a student?

Noam and Alvin. Sharing a common interest with your advisor holds significant importance. Initially, one might not have a clearly defined research question, but through ongoing discussions and shared interests, a focused and refined question can emerge. This collaborative approach facilitates a deeper exploration of topics and fosters a productive mentorship dynamic, guiding the research process towards meaningful outcomes.

How often do you talk to your students?

Alvin. I schedule two coffee meetings each week in the morning, valuing these as a foundation for intellectual development. This approach is based on the belief that engaging conversations are a crucial source of intellectual insights, fostering an environment where innovative ideas and solutions can emerge from collaborative discussions.

How do you approach learning and developing your expertise to effectively contribute to solving problems?

Noam. I approach my work in two distinct methods. Initially, I dive deeply into the theoretical aspects of mechanism design, starting with foundational textbooks such as the Mas-Colell book to grasp the underlying concepts thoroughly. This

“original mode” focuses on building a strong theoretical foundation. Conversely, the second approach emphasizes practical application: setting aside theory to experiment with solutions directly in real-world scenarios. This mode prioritizes action and experimentation to see how concepts perform outside of theoretical frameworks, fostering a balance between theoretical understanding and practical applicability.

Alvin. At all times, delving into books and consulting with field experts, such as doctors for the kidney exchange problem, is essential. This approach of dynamic engagement and exchanging ideas with specialists enhances the depth of understanding and brings new perspectives to the forefront. It underscores the importance of interdisciplinary communication in refining research questions.

As a final insight, could you provide an elaboration on a real-life problem in the field of mechanism design?

Noam. During my tenure with Google Ads, I observed that despite advertisers having budgets, they often opted for simple models over complex ones. This realization highlighted a disconnect between the sophisticated models developed in theory and the practical needs of advertisers. The simplicity and applicability of models, rather than their theoretical sophistication, were more aligned with the advertisers’ objectives, showing the importance of developing solutions that are both effective and user-friendly for real-world applications.

Alvin. Over time, the dynamics within the field of kidney exchange have evolved significantly. Initially, the focus was on individual practitioners and patients. However, as the system has grown and developed, the involvement of hospitals and kidney exchange program directors has introduced strategic considerations into the process. This shift reflects a broader and more complex approach to organizing kidney exchanges, highlighting the need for strategic planning and coordination at higher organizational levels.

2.1 Further Q&A with Noam Nisan

After the meeting, Noam was kind enough to share some additional thoughts ranging from his own research interests to general directions within the EconCS sub-community, and other advice for researchers in the space.



What can machine learning and human behavior tell us about each other? And what other sorts of cross-pollination (e.g. human learning? machine behavior?) should we be thinking about?

For a long time, we’ve thought that machine learning can teach us about human learning—that by uncovering or recreating the cleverness of human reasoning by developing algorithms and mechanisms, we can understand human intelligence more. I agree with this, but with a twist: I think machine learning helps us understand how much *less* spectacularly we should think of ourselves.

Sometimes when people work with ChatGPT and get annoyed by foolish errors, what really strikes me is that these errors are exactly the same kinds of things that

humans do. One of my favorite examples: if you ask ChatGPT to solve a simple math problem that it's not smart enough to do, it'll give you an answer where every line looks like it's trying to develop something sort of reasonable, but if you really look at the whole thing it's complete nonsense. And it always reminds me exactly of the exams of students that don't know the material, but you can see that they're trying their best to write things that sort of seem correct to get partial credit. So I think that to a very large extent, one of the things we can learn from ML is how humans behave in ways that are *not* so clever.

We humans think of ourselves as being very clever, say by playing chess, but now that we have machines that can play chess perhaps we see that chess does not represent the highest level of intelligence. Maybe object recognition is a higher level of intelligence—and once we have machines that can do that, we'll move onto solving math problems, or cleaning a room, or some other task entirely. But every time we get a machine that can do something new, one way to look at it is as demystifying what humans do, and maybe thinking of ourselves a little less highly.

So would you say that there's this overall perception that humans are somehow “better” than AI and machines, but you're trying to challenge this notion?

First of all, if you're not assuming some sort of God or dualism or other principle, humans are also biological, psychological machines, right? So the question is: will we be able, in the near future, to build machines that are better than us in various tasks? Nobody can know for sure, but I don't see any reason why not, sooner or later. And I think that's fine—maybe at a certain point, we'll just have to accept that our worth as human beings is not because of our intelligence; maybe it's our morality, or decision-making, or something else, but not just pure intelligence.

So that's a philosophical point of view. But practically speaking, yes, I do believe that computers will eventually outperform humans. Right now everyone's talking about LLMs; maybe next year we'll all be using something else entirely. I don't see any major reason to believe that in 5, or 10, or 20 years, machines won't be better than humans at proving mathematical theorems—like now, there's no human that can outrun a car, and so be it. Maybe our generation was lucky that we could find gainful improvement by proving theorems, and that may not be the case 30 years from now.

That's definitely a fascinating perspective on the future of AI. To take it back to the (recent) past, you've also done some work on blockchains and cryptocurrencies. Do you view this line of work as entirely separate from your work on behavioral models, or is there some insight that might be gained at the intersection of the two?

It's still disjoint, but related in a sense. Blockchains are a technical means to the goal of finding a mechanism that works without a centralized authority, and there are various computer science, economic, and engineering challenges in building them, but their overarching reason—what's really interesting about them—is the social visions that they allow. This could be similar to the introduction of the Limited Liability Company a few hundred years ago, which represented a different

form for humans to organize themselves in a way that turned out to be extremely socially important. Maybe blockchains will bring about another type of revolution into self-organizing, self-operating, complex networks. From this point of view, I'm interested in the social aspects of these networks, and of course social aspects are related to human behavior.

But I don't think we've reached a level of understanding of human aspects of behavior to then understand the social implications yet. Right now, we're focused on the technical questions of how to build blockchains, and are basically operating off intuitive assumptions of how people behave. We don't have nice theory or empirical results yet. But I believe this is definitely a direction that science should advance towards—at the end of the day, social structures depend on the behavior of the humans inside.

What's one problem within this space that you would love to see more researchers from the EC subcommunity work on?

I would be interested in people trying to develop the theory of tokenomics: an economic theory of these types of blockchain networks, and how they should behave. A lot of research has been done on the details of how all these blockchains differ in their abstractions and implementations, how you can maybe cheat a bit even with a small minority of faulty nodes, how you can get more computational efficiency, and so on. But I'm more interested in supposing that we've reached a situation where these decentralized networks work—now, what do they do? How do we understand their dynamics? There's economic and design questions like token pricing, voting, fees, and other system constraints. In general, these are complex systems, and I would like to see more work that tries to formalize and capture and understand the high-level aspects of the operation of these systems, under the abstraction that they work as intended.

And what would you say to the blockchain skeptics out there—any interesting non-blockchain behavioral questions you have in mind?

First, they're very right to be skeptical. I think it's a very big question whether blockchains will succeed—I would not be at all surprised if everything completely fails tomorrow, nor would I be surprised if in 10 years from now blockchains run half the world. So I really don't know what will happen—there are many possibilities that we could see.

With regard to general behavioral questions, this is an area that I believe is very important, but it's challenging for me to highlight specific questions. My tendency is generally to try to simplify things until I get a question that I can handle, but with humans, sometimes you can oversimplify and lose essence. This makes behavioral questions very tricky, and I have to admit, less aligned with my research approach. I hope that people who have this touch of combining psychology, computer science, and economics do more work in this area, but I really am not able to suggest concrete directions yet.

In general, how do you find balance between trying to gain expertise in a single focused area versus staying open to exploring a broad range of new areas (like behavioral modeling)?

That's a good question, especially as computer science is rapidly growing and changing. Back when I was a PhD student in the 80's, you could basically know all of theoretical computer science—I would say a senior grad student could feasibly go to STOC or FOCS and understand maybe 80% of the papers there.

That's not the case today anymore. It was certainly a fun time to grow up scientifically—I tried to go deep in one area and wide in general, because that's the standard recommendation. But I'm afraid that today it's difficult to be wide in general—so how wide do you go?

It's difficult to say, but you do need to constrain yourself to *something*, after all. I think you should try to find topics that will be important for the future on one hand, and on the other hand where we can make some meaningful progress today. Think about building a time machine: perhaps one of the most important questions for the future, but we currently have no idea how to do anything useful towards this direction. Easier said than done of course, but I think every scientist, every generation, faces this question.

So what's next for you and your research agenda?

I really like the blockchain world, and I'm trying to understand more stuff there. The nice thing about being a tenured professor is that understanding doesn't necessarily have to be the same thing as writing papers at this point—for me, I really like going to talks to get the basic ideas, and then if I'm interested in completing the picture, I try to figure things out myself by reading the details of the paper.

Another thing I'm interested in is how machine learning relates to algorithmic game theory, like in marketplace dynamics with an ML component—how does regret minimization work in markets? how do autobidders operate?—and other such questions combining AGT and markets with learning.

Finally, what's one interesting, non-CS fact about you?

Well, of course one never wants to share the *most* interesting answers to this question... but I suppose I could let some people here know that I have been doing Israeli folk dance for over 40 years!

3. SPOTLIGHT TALKS AND INTERVIEWS

3.1 Modibo Camara - Computationally Tractable Choice



In the first spotlight talk, **Modibo Camara** (Stanford GSB) presented a paper that studies choice when agents face computational constraints [Camara 2022]. Agents have only a limited amount of time to make their choices, and making good decisions takes time. In particular, always making the optimal choice may be too time consuming. For these reasons, agents often rely on behavioral heuristics to make their choices. An important example is choice bracketing—partitioning choices, and then optimizing within each element of the partition.

Are such behavioral heuristics a necessary response to computational constraints? Modibo gives a two-folded affirmative answer to this question, using concepts from theoretical computer science to define tractability. A choice rule is *tractable* if it can be implemented in polynomial time by some algorithm. First, if a choice rule that is consistent with the expected utility axioms is tractable, then it is observationally equivalent to some form of choice bracketing. Second, an expected utility maximizer who is constrained to use tractable choice rules can be better off by using heuristics that would not appear rational to an outside observer.

More precisely, consider an agent who has to specify a choice rule for a set of menus, which satisfies some richness conditions. Each menu consists of lotteries over eventually-zero sequences of rational numbers. A product menu is the Cartesian product of its *partial menus*, defined as the (marginal) lotteries that the agent may choose for each entry of the sequence.

Within this framework, rational and tractable choice rules that satisfy a symmetry condition can be rationalized by additively separable Bernoulli utility functions. This implies that, for product menus, such choice rules are observationally equivalent to narrow choice bracketing—that is, optimizing over partial menus separately. More generally, a choice rule is rational and tractable if and only if it is observationally equivalent to *dynamic choice bracketing*, which allows the relevant brackets to change in the process of making the choice.

For the final result, consider an expected utility maximizing agent who can only choose tractable choice rules. Surprisingly, her Bernoulli utility function may be such that the following is true: There is no rationalizable tractable choice rule that guarantees the agent at least one-half of a menu's optimal value, while there are non-rationalizable tractable choice rules which do. This result has a particularly interesting takeaway: Once we accept the idea that real-world agents have limited computational capacity, behavioral heuristics that may appear irrational from the perspective of an outside observer can actually be more effective than fully rationalizable choice rules.

3.1.1 Interview with Modibo Camara. In an interview after the Winter meeting, Modibo Camara talked about the process that led to his paper and his perspective on future research in the field.

How did you get interested in the theory of simplicity, and more specif-

ically how did you come up with the idea for your paper?

When I started my Ph.D., my initial impression was that economic models often assume that people can do things that are too complicated. As a consequence, some of the recommendations of these models felt unrealistic and not very applicable. This made me think that taking complexity seriously may be of great importance for economic theory.

The research process that eventually led to my job market paper started during my first year. I was not fully satisfied with the axiomatic approach to decision theory: I believed that, if complexity were to be taken into account, then an agent may be better off by choosing simpler behaviors, that—despite violating the consistency requirements of expected utility theory—are approximately optimal. In particular, I had a feeling that when considering large numbers of decisions, a complexity constrained agent had to rely on some sort of narrow bracketing. I formalized this idea with my paper.

What other topic are you working on that you find the most exciting?

Related to bounded rationality and computer science, one topic that a collaborator and I have been thinking about recently is Artificial Intelligence (AI) alignment. One way to think about AI is as a tool to overcome humans' bounded rationality. From that perspective, if humans did not find computation taxing, there would be no need for AI. But there is still a computational bottleneck. In order to delegate a decision to the AI, the human needs to communicate her preferences to the AI, either directly or indirectly (e.g., through past choices). Studying the complexity of such communication seems important for identifying settings in which AI can be effective, as well as settings in which AI may never be fully aligned with human preferences.

More broadly, what topics at the intersection of computer science and bounded rationality in economic theory do you think are most promising for future research?

One interesting topic for future research is recommender systems. These are algorithms used by many online platforms, and their purpose is to resolve market failures that arise because users are boundedly rational. Take Netflix, for example. One of the main services that Netflix provides is that it uses behavioral data to filter a massive catalog of movies and television series into personalized recommendations for its users. Without this algorithm, users that struggle to process Netflix's massive catalog would likely end up making significantly worse choices. I think recommender systems represent a promising application area for work on bounded rationality, since they drive a huge amount of consumption online and are fundamentally a response to users' cognitive limitations. Optimistically, the EconCS community should be able to explain why these systems exist and provide guidance on how to design them.

3.2 Ryan Oprea - What Makes a Rule Complex

In the second spotlight talk, **Ryan Oprea** (University of California, Santa Barbara) presented experimental work which investigates what makes rules complex for human beings [Oprea 2020].

The paper aims to make a first step into ranking decision tasks and predicting which sub-optimal but less costly algorithms, such as heuristics, will be favored by humans. To this end, it considers the following experimental setting to measure subjective cognitive costs of performing algorithms directly and understanding what makes a rule complex for humans. Agents are asked to implement some rule—a sequence of choices as a function of a sequence of events. For example, an agent may be asked: Pick x until you see a , then switch to y . They get paid if and only if the task is performed correctly. Then, the experimenters elicit their willingness to pay to avoid to be asked to implement it again in the future. Which rules are more complex? Why?

In order to answer these questions, rules need to be systematically comparable. To this end, the paper represents rules as automata. That is, each rule is decomposed into a set of states and transitions: Each state is identified with an action to take, and is graphically represented as a circle. Each transition is identified with an event that triggers a shift from a state to another, and is graphically represented as a connecting arrow.

This approach reveals several recurring patterns of complexity. While agents make few mistakes, they have high willingness to pay to avoid to implement the proposed rules. Moreover, agents' willingness to pay varies significantly with the number of states of a given rule. Consistently with intuition, the data show that costs are increasing in the number of states and transitions. Interestingly, agents do not efficiently represent algorithms to themselves: In particular, they fail to see that some states and transitions may be redundant.

The above results notwithstanding, finite automata may still miss important aspects of complexity. Indeed, more sophisticated automata (known as pushdown automata) can better fit the data, and even richer languages—such as Turing machines—may be needed to fully understand the costs of algorithms.

In closing, Ryan stressed how we only just started to understand how to design simpler social rules that can work better in practice. In particular, this work is a call for future research to build models of bounded rationality rooted on empirically grounded characterizations of complexity.

3.3 Gali Noti - Learning When to Advise Human Decision Makers

In the third spotlight talk, **Gali Noti** (Cornell University) presented joint work with Yiling Chen on learning when to advise human decision makers [Noti and Chen 2023]. AI-assisted human decision-making, in which a machine learning algorithm is intended to *help* rather than *replace* a human decision-maker by providing advice (in the form of predicted *risk scores*), is becoming an increasingly widespread paradigm in fields ranging from criminal justice to healthcare. Gali's work asks the following question: to best improve human decision-making, *when* should algorithms provide advice?


A natural guess is that algorithms should always provide advice—and indeed, this is generally what happens in current practice—but it turns out that this may not actually be the case. In fact, experience suggests that humans and algorithms have complementary strengths—algorithms are not perfect, and there are still areas where humans can outperform algorithms. Another empirically observed phenomenon is a *scarcity effect*, in which humans respond more strongly when advice is provided

less frequently. Taken together, these facts point towards a *responsive advising approach*, where the AI only provides the advice when it is expected to improve the decision. Is such an approach possible? And if so, does it actually help?

Using the context of predicting recidivism in pre-trial release decisions, this work considers an algorithmic assistant consisting of a risk assessment algorithm (produces a risk score) and an advising policy (decides whether or not to provide advice). The results? The algorithmic assistant learned to advise with 74.1% accuracy, and the assisted human decisions were approximately as accurate as the algorithmic predictions. In short—yes, responsive advising is both possible *and* helpful!

Other empirical findings included an illustration of the scarcity effect, as well as evidence of “learning on both sides” (by both the human and the algorithm). Finally, Gali emphasized that many richer forms of human-AI interaction remain unexplored. Moving forward, as we continue to navigate the various tradeoffs between learning from algorithms and preserving human strengths, it will become increasingly important to understand how humans and algorithms behave in collaboration with each other.

3.4 Nicole Immorlica - Data, the Fundamental Particle of Interaction



The final spotlight talk of the day was given by **Nicole Immorlica** (Microsoft Research) on data, the fundamental particle of interaction. Nicole began with the motivating observation that one of the most basic assumptions underlying much econCS research—that rational agents make decisions by maximizing expected utility according to prior distributions—is *not* actually an accurate model of human behavior. In reality, priors are difficult to formulate precisely; instead, humans tend to operate off past

experiences. For large corporations and other organizations, this is often in the form of massive datasets; on the level of individual human decision-making, this is usually based on *anecdotal experiences*.

Anecdotes are a useful tool for representing non-experts (who lack sophisticated formal models of the world), and considerations such as data privacy or other constraints. In a stylized model, the state of the world is parametrized by some value unknown to the agent, and an *anecdote* is a noisy representation of the state of the world. Now, an agent receives anecdotes one at a time, and can choose to remember or forget each one. How many anecdotes does an agent need to use to form a “good” estimate of the state of the world?

If the agent has perfect memory, the Central Limit Theorem implies that the empirical mean of T anecdotes has loss $O(1/T)$. But what if the agent only has bounded memory? The simple but naive strategy of remembering only the most recent anecdotes requires $O(T)$ memory to match this loss bound of $O(1/T)$. In the busy and chaotic world we live in today, we might hope to ask: can we get away with remembering any less?

Good news for the forgetful folks: yes! It turns out a form of *strategic forgetting* can achieve this same $O(1/T)$ loss with memory only logarithmic in T . The complete specification and analysis of the algorithm are fairly technical, but also enlightening and exciting. One key insight is the notion of setting the optimal “step

size” parameter—too small, and the empirical distribution barely changes (reflecting an agent who stubbornly refuses to learn from experience); but too large, and the empirical distribution fails to converge (reflecting an agent who is too easily swayed and keeps changing her mind). This latter observation in particular suggests a rationale for *confirmation bias*, the widespread behavioral phenomenon in which an anecdote must be at least somewhat close to an agent’s current beliefs in order for her to remember it. This has interesting implications—oftentimes we think of “rational” and “unbiased” as interchangeable terms, but in this case (and perhaps in others), bias may be rational after all.

To close, Nicole called for researchers to begin investigating a broader agenda of viewing *data* as *anecdotes* correlated with the state of the world. This framework can be extended beyond this talk’s decision theory problem, to other contexts ranging from communication and persuasion games to mechanism and information design. More generally, it’s time to start viewing data as the *fundamental particle* underlying much of human behavior and decision-making, and develop economic models that reflect this new paradigm shift.

3.4.1 Interview with Nicole Immorlica. In an interview after the meeting, Nicole gave further insight on modeling data as anecdotes, and shared some broader perspectives on behavioral research in the growing field of economics and computation.

What inspired you to start reasoning about models using data and anecdotes instead of distributions? Was it pure intellectual curiosity, motivated by a product team, or something else entirely?

I think it was mostly intellectual curiosity, but this curiosity was initiated by observations from product teams. My first paper in this space (*Incentivizing Exploration with Selective Data Disclosure*, joint work with Alex Slivkins [Immorlica et al. 2018]) was motivated by platforms like Yelp—we wanted to think about how Yelp uses reviews (which are anecdotes or data points) to persuade people to try different restaurants. Why is this a natural model? Why does Yelp actually aggregate past reviews instead of just making up arbitrary scores? Perhaps it’s because that’s how humans reason—we look at experiences others have had in the past. So this was the first moment where I thought this kind of framing makes sense—Yelp is using reviews because humans reason better about concrete data than they do about abstract messages.

My next paper in this framework, *Communicating with Anecdotes* [Haghtalab et al. 2022], took this further—if people are in the business of sharing facts and data points rather than beliefs and messages, how does conversation work? What happens if we can’t just say anything we want (the typical *cheap talk* assumption underlying much of the economics literature), but can only communicate things that have actually happened? In this world, you can think of this as only being able to communicate *true news* via facts, evidence, and anecdotes. What beliefs then get induced? This paper was also influenced by many fascinating conversations I had with Nancy Baym, a communications scholar at Microsoft, about why people communicate, as well as some psychology literature that one of my coauthors, Markus Mobius, was reading at the time. It was fantastic to have access to people outside the EC subcommunity who think about things from different perspectives.

This paper I presented at the winter meeting was my third in this line of work. It combines the previous idea that we think about data and anecdotes with this concept of limited memory that Markus was intrigued by from a book he was reading (*Foundations of Human Memory* by Michael Jacob Kahana [Kahana 2012]).

Zooming into some technical details of your talk: Can we say anything constructive about distributions that are not necessarily single-peaked? (Using your coffee example from your talk, what if there's a "good barista" and a "bad barista" so the distribution of coffee quality is bimodal?)

All of the papers I discussed have some kind of single-peaked assumption; this serves as a focal point to narrow in on, so that the mass of the distribution is where the mean is. If you have a bimodal distribution, instead of converging, gradient descent might just keep jumping around between the two local maxima. It also depends on what exactly you're trying to learn in this case—the average quality of the coffee shop, or the quality of one barista versus the other? In the latter case, I actually do have two single-peaked distributions, just obfuscated (and if I can tell which samples are from which distribution, then we can apply the multi-dimensional extension of our algorithm).

But in general, you do need some kind of structure on the distribution—think of it as coming from the assumption that there's a ground truth of the state of the world, and observations are noisy perturbations. If the noise is zero mean, then you'll end up with a single-peaked distribution.

Have you thought about modeling other behavioral phenomena, such as rose-colored glasses, memories drifting over time, or more strongly remembering extreme memories?

For sure, I'm quite interested in looking at how cognitive constraints may lead to behavioral phenomena. I would love to assume that people are trying to optimize their true underlying utility but have limited abilities, and when they try to find the optimal algorithm subject to these constraints, it turns out to have behavioral implications. It would be cool to come up with model where the human is trying to optimize subject to being able to run an algorithm in their head, and as a feature of the best such algorithm, ends up exhibiting a particular behavior (in the talk I presented, confirmation bias arose from humans having limited memory).

Remembering extreme events is something I've talked about a bit with Nageeb Ali. You could imagine a setting where the goal is to learn the best action out of several choices (rather than a mean estimation problem), so you would love to run a multi-armed bandit algorithm, but you can't quite do so because you only have anecdotes and maybe also limited memory. What could you do instead? One thing you might want to do is to mimic UCB by remembering the best thing you've seen so far. Maybe if you have enough memory to remember something about every arm, you'd also want to remember the really terrible points so you never try them again. So interestingly, both the phenomena of remembering extreme events, or just the very positive ones, could arise from trying to mimic UCB.

You hinted at other contexts where this sort of paradigm shift from

distributions to data could be applied, like in game theory and mechanism design. What's one direction that you find particularly exciting?

I'm starting to try to understand the impact of LLMs on how we write our economic models and games. Here, I think of LLMs as both content creators and content aggregators. In the former aspect, there's this whole theme of generating the content (data points) that we see. How does the existence of these AI machines that can contribute to this environment influence what data we create? If the LLM is creating data and also competing with me for attention, does this influence how I create data myself?

Regarding the other concept as an aggregator, so far I've been operating under the basic assumption that I can see and communicate and remember a data point. But with a summarization tool, I can remember *combinations* of data points instead. If I can remember arbitrary combinations, then I get back to standard models of distributions, but what if there are constraints on the combinations? Further, what if the LLM wants to influence humans to think a certain way? On one hand, LLMs can support human learning by providing combinations of data points; on the other hand, these combinations may be biased by the LLM. More broadly, how do LLMs affect how people learn (especially if incentives are misaligned)?

More generally, what is your process for coming up with interesting research problems?

I think everyone has their own process, and what inspires people may be different. For me, I'm much more into the modeling aspect than the theorem-proving aspect, so I try to look for something I don't really understand and figure out why it's happening. In this instance: people often make anomalous, seemingly sub-optimal decisions. Why? Maybe it's because they have limited memory. Overall, my research is a tool for helping me understand the world around me.

There's a big gap between figuring out what you're interested in the world and what questions to ask. Another thing I've been interested in for a while but haven't formed the right questions around is the concept of identity. What is identity? How does it matter to me that I identify as a woman, as both an economist and a computer scientist, as a Microsoft employee, as a member of my family? How do identities change? How do conflicts between my perception of my own identity and others' perceptions of me influence my actions? There's this quote from a Walt Whitman poem: "Do I contradict myself? Very well then I contradict myself, (I am large, I contain multitudes.)" This is another hint that we do things that may seem absurd from a rationality perspective, but maybe that's because we are not single rational actors! Even a single individual is composed of many forces pulling in different directions—how can we use this notion of identity to think about how we act as a result?

So to get ideas, I read poetry. I read related works (Rachel Kranton and George Akerlof have some great stuff). I talk to people outside of the subfield. Then I start trying to write down models.

What's one random fun fact about you?

I really value breadth of experience, both academically and personally. For example, just pre-pandemic I took a poetry course at MIT because I'm interested in

engaging with fields beyond STEM. In fact, one of my favorite classes in undergrad was German film post-World War II!

REFERENCES

- CAMARA, M. K. 2022. Computationally tractable choice. In *Proceedings of the 23rd ACM Conference on Economics and Computation*. EC '22. Association for Computing Machinery, New York, NY, USA, 28.
- DREYFUSS, B., GLICKSOHN, O., HEFFETZ, O., AND ROMM, A. 2022. Deferred acceptance with news utility. Tech. rep., National Bureau of Economic Research.
- GONCZAROWSKI, Y. A., HEFFETZ, O., AND THOMAS, C. 2023. Strategyproofness-exposing mechanism descriptions. Tech. rep., National Bureau of Economic Research.
- HAGHTALAB, N., IMMORLICA, N., LUCIER, B., MOBIUS, M., AND MOHAN, D. 2022. Communicating with anecdotes. *arXiv preprint arXiv:2205.13461*.
- IMMORLICA, N., MAO, J., SLIVKINS, A., AND WU, Z. S. 2018. Incentivizing exploration with selective data disclosure. *arXiv preprint arXiv:1811.06026*.
- KAHANA, M. J. 2012. *Foundations of human memory*. OUP USA.
- KLEINBERG, J., MULLAINATHAN, S., AND RAGHAVAN, M. 2023. The challenge of understanding what users want: Inconsistent preferences and engagement optimization. *Management Science* 0, 0, null.
- KLEINBERG, J., OREN, S., AND RAGHAVAN, M. 2016. Planning problems for sophisticated agents with present bias. In *Proceedings of the 2016 ACM Conference on Economics and Computation*. EC '16. Association for Computing Machinery, New York, NY, USA, 343–360.
- LI, S. 2017. Obviously strategy-proof mechanisms. *American Economic Review* 107, 11, 3257–3287.
- MCILROY-YOUNG, R., SEN, S., KLEINBERG, J., AND ANDERSON, A. 2020. Aligning superhuman ai with human behavior: Chess as a model system. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. KDD '20. Association for Computing Machinery, New York, NY, USA, 1677–1687.
- MEISNER, V. 2023. Report-dependent utility and strategy-proofness. *Management Science* 69, 5, 2733–2745.
- NOTI, G. AND CHEN, Y. 2023. Learning when to advise human decision makers. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI-23*, E. Elkind, Ed. International Joint Conferences on Artificial Intelligence Organization, 3038–3048. Main Track.
- OPREA, R. 2020. What makes a rule complex? *American Economic Review* 110, 12 (December), 3913–51.

Decision Theory, Subjective Uncertainty, and Computer Science

GIACOMO LANZANI

Harvard University

I argue that further integration between Decision Theory and the methods of quantifying complexity and evaluating performance in Computer Science is valuable. I review [Lanzani 2024] as an illustration of this combination.

Categories and Subject Descriptors: J.4 [Social and Behavioral Sciences]: Economics

General Terms: Economics, Performance, Theory

Additional Key Words and Phrases: Learning, Misspecification, Ambiguity, Convergence

1. OBJECTIVELY NORMATIVE DECISION THEORY

Decision Theory, as a subliteration within Economics, has always faced a tension between a normative interpretation and a descriptive one, between whether the decision criteria and axioms proposed should be interpreted as desiderata or postulates about the actual behavior of economic agents. One area in which the normative interpretation has been prevalent is decision theory under uncertainty.

However, for Subjective Expected Utility (SEU) and Bayesianism as postulated by [Savage 1954] (see [Cerreia-Vioglio et al. 2013] for its “belief over models” interpretation pursued here), the normative justification is an internal, or subjective, one. Seeing the axioms corresponding to the decision criterion, the decision maker feels compelled to adhere to them. Still, they do not imply that a decision-maker following them will perform particularly well in an objective environment. In static decision problems, there is no reason to expect such a good objective performance, as the decision maker is not omniscient. However, in dynamic environments, the gap between subjectively and objectively good decisions is elegantly closed for Bayesianism by Bernstein von Mises type theorems (cf. [Doob 1949; Breiman et al. 1964]), which show that if a Bayesian statistician is correctly specified, and the set of models considered is well-behaved, they will eventually concentrate their beliefs on the true data-generating process (DGP). Consequently, a correctly specified Bayesian agent will eventually take the ex-post optimal course of action.

In subsequent years, several different decision criteria have been proposed and axiomatized for decision theory under uncertainty. Many of these decision criteria have been motivated with an internally normative motivation; see the material reviewed in [Gilboa 2009; Gilboa and Marinacci 2016; Hansen and Marinacci 2016]. Chief among the normative motivations for the departure from subjective expected utility was a concern for complexity: many decision problems faced by economic agents are so complicated that it is impossible to quantify them with a set of proba-

Author’s address: giacomolanzani34@gmail.com. I thank Roberto Corrao, Federico Echenique, Drew Fudenberg, and Irene Lo for their helpful comments.

bilistic models that could be fully trusted. Concretely, this impossibility comes from the fact that a set of relatively simple (e.g., parametric) models will be very prone to misspecification, i.e., not to include the true DGP. In contrast, an excessively large set of models will make Bayesian updating inconsistent (cf. [Diaconis and Freedman 1986]), or less dramatically, it will make convergence excessively slow.

However, so far, less progress has been made on whether these decision criteria *objectively* guard against the problems that have motivated their introduction. This question can be asked for countless decision criteria under uncertainty, e.g., the maxmin model [Gilboa and Schmeidler 1989], Choquet Expected Utility [Schmeidler 1989], multiplier preferences [Hansen and Sargent 2001; Strzalecki 2011], variational preferences [Maccheroni et al. 2006], models that combine worst and best case scenarios [Ghirardato et al. 2004; Gul and Pesendorfer 2014] or even general uncertainty averse preferences [Cerreia-Vioglio et al. 2011].¹ Do these criteria perform well in an objective situation in which a single probability measure would have been misspecified? How do we even operationalize the request for good performance? Does the reduction to a smaller set of DGPs paired with caution in using them overcome the inconsistency (or slow consistency) problems faced by the Bayesian paradigm with an extremely large set of models? I next turn to the discussion of why the approaches developed in Computer Science could be useful in formalizing and answering these questions.

2. WHY COMPUTER SCIENCE CAN HELP

The Computer Science literature has a long tradition in three areas suited to interact well with the goal of providing an objective assessment of the performance of decision rules under uncertainty: i) Criteria for the *objective* evaluation of decisions in repeated decision problems; ii) An emphasis on the *speed* of convergence; iii) An emphasis on the *complexity* of algorithms and decision rules.

Objective Evaluation Criteria. First, the Computer Science literature can contribute by providing a variety of criteria used to evaluate the performance of decision rules. The most studied is the concept of no-regret dynamics (see, e.g., [Noam Nisan 2011]). This requirement proved particularly useful for game-theoretic settings. Still, the dual representation as a game against Nature featured by uncertainty-averse decision criteria, [Cerreia-Vioglio et al. 2011] suggests that such techniques could prove useful even when a single decision maker is involved.

A less widely-used objective valuation criterion that suits the problems faced in Decision Theory is the advice-augmented dynamic performance. Loosely speaking,

¹Of course, a dynamic version of these criteria has been formulated and studied, and the descriptive predictions for a decision-maker using such criteria and learning have been explored (see, [Marinacci 2002], [Epstein and Schneider 2007], and [Battigalli et al. 2019] as well as the extended references in [Hansen and Sargent 2011] and [Ilut and Schneider 2022] for results with particular attention to their macroeconomic and finance applications). The emphasis of this article is on decision theory under uncertainty, but I believe a similar perspective can also be useful in evaluating whether theories of choice under risk built on the promise to protect from uncertainty in the evaluation as [Cerreia-Vioglio et al. 2015] or correlation (see, e.g., [Lanzani 2022] for a unifying treatment) objectively deliver such protections.

it considers settings in which the decision-maker has access to the repeated predictions of an algorithm or a model and requires that, as long as the algorithm is not corrupted or the model is correct, the DM obtains the optimal performance if they completely trusted the advice. This approach has been successfully applied in many settings (see, e.g., [Mahdian et al. 2012; Banerjee et al. 2022]). This evaluation criterion combines particularly well with the economic approach in that economic agents are often endowed, even under non-SEU decision criteria, with a set of (possibly simplified) models they use to inform decisions.

Speed of Convergence. Second, the Computer Science literature pays special attention to the quantitative determination of the performance of a dynamic strategy. No regret, especially in individual problems, is hardly a goal for computer science but rather a given. The emphasis is more on determining the speed at which the regret decays to 0 (see, e.g., [Cesa-Bianchi and Lugosi 2006]).

Again, this focus seems to fit well with Decision Theory under uncertainty, especially given its motivation; see, e.g., this passage from the survey in [Gilboa 2023]:

“... my personal view is that it is wrong to assume that the only rational way to make decisions is to adopt a prior and follow SEUT ... If there is an infinite horizon of learning periods ahead of us, the choice of such a prior may be almost immaterial: as long as the prior is sufficiently open-minded, the underlying process would be learnt. But there are too many problems, ranging from wars to climate change, where we simply don’t have the time to learn the underlying process ... In these cases, I believe that it may be more rational to admit that we do not know the probabilities than to pretend that we do.”

But then, if non-SEU decision criteria are proposed as responses to excessively slow learning, their dynamically accumulated loss should be quantified and compared with the speed of learning with a large SEU model, a task for which the Computer Science tools seem particularly apt. These trade-offs are also at the core of the machine learning literature that proposed notions of the complexity of models (like Vapnik–Chervonenkis dimension and Rademacher complexity) that connect the speed of convergence to model complexity.²

Complexity. Third, uncertainty-averse decision criteria are often defended with a complexity argument. The idea is that postulating a comprehensive, all-encompassing model of the economic environment and performing dynamic optimization and learning within such a large model is computationally infeasible. Therefore, it is informally claimed we should use simpler scenario(s) but cautiously use them. However, these cautious criteria are often highly nonlinear, and thus, the corresponding optimization processes become more involved.

Although many papers illustrate the tractability of these problems in specific settings or under particular parametric assumptions, almost no work deals explicitly with the tradeoff between the complexity of Bayesian updating in a large SEU-Bayesian model versus the complexity in the optimization problems that depart

²The study of these tradeoffs dates at least back to [Valiant 1984], although most of the emphasis has been on avoiding the problem of overfitting.

from SEU. Here, computer science seems again to be able to provide relevant insights, as the quantification of complexity has always been a central topic. Indeed, recent years have seen contributions in this direction, see [Echenique et al. 2011; Fudenberg et al. 2022; Camara 2022].

3. ILLUSTRATION: [LANZANI 2024]

This section briefly summarizes the findings of [Lanzani 2024]. There, I take a class of *static* decision criteria under uncertainty motivated by robustness to misspecification [Hansen and Sargent 2001; Cerreia-Vioglio et al. 2022]. I then combine it with some objective evaluation criteria borrowed from the computer science literature, specifically the requirement of a minmax guarantee and an adaptation of the “advice-augmented” consistency mentioned in Section 2. I show that this requirement pins down a unique *dynamic* way to adjust those decision criteria in the face of accumulated evidence. I leverage this normative result to characterize the long-run behavior we should expect when agents employ those dynamically adjusted decision rules and when they depart from it in the direction of being excessively demanding or lenient in evaluating their model performance. The subsequent summary borrows extensively from [Lanzani 2023].

I consider an agent that repeatedly chooses among actions $a \in A$ that induce an unknown distribution over outcomes $y \in Y$. An objective data-generating process maps current actions into a probability distribution over outcomes $(p_a^*)_{a \in A}$. The agent does not know p^* but rather envisions a simpler set of models Q , where each $q = (q_a)_{a \in A} \in Q$ is also an action-contingent collection of probability distributions over outcomes. The agent has a utility function u over the joint realization of actions and outcomes. At every period, this choice is determined by maximizing the average (with respect to a belief $\mu \in \Delta(Q)$) of robust control assessments, where each assessment uses a different structured model as the benchmark:

$$\int_Q \min_{p_a \in \Delta(Y)} \left(\mathbb{E}_{p_a} [u(a, y)] + \frac{1}{\lambda} R(p_a || q_a) \right) d\mu(q). \quad (1)$$

Here, R denotes the relative entropy, λ is a real number, and the decision maker trades off between the performance of the action under the conjectured models in Q and worst-case scenarios that are not too far in terms of relative entropy.

I introduce endogeneity in the misspecification concern: the better the structured models explain the past, the less concerned the agent is. I first establish a normative result. If the agent wants, across a large set of environments, to be guaranteed to both:

- (1) Always achieve the minmax payoff;
- (2) Achieve the (approximately) ex-post optimal payoff if their model is (approximately) correct,

then, they should evaluate their model using a log-likelihood ratio and keep the level of concern for misspecification proportional to this log-likelihood ratio. I also show a partial converse, in that every rule of adjustment of the concern for misspecification that is either asymptotically faster or slower than this fails one of these two properties. Observe that requirement (2) can be read as an approximate

version of the advice-augmented dynamic performance criterion, where predictive algorithms are replaced by predictive economic models.

I then move to the descriptive analysis, where I allow departures from this normative benchmark and consider agents that are too demanding in evaluating the models' performance. Similarly, I allow the opposite case in which the agent is too lenient in evaluating their model and attributes too much unexplained evidence to sampling variability.

With this, I characterize the possible long-run behavior of these different types of misspecified agents. The limit actions of the lenient type must converge to a Berk-Nash [Esponda and Pouzo 2016; ?] equilibrium, i.e., to an SEU best reply to beliefs supported on the models closest to the data-generating process. In mathematical terms, the limit action a must satisfy the following fixed point condition:

$$a \in \operatorname{argmax}_{a' \in A} \int_Q \mathbb{E}_{q_{a'}} [u(a', y)] d\nu(q)$$

with

$$\operatorname{supp} \nu \subseteq \operatorname{argmin}_{q \in Q} R(p_a^* || q_a).$$

Instead, overemphasis on the model's failures in explaining the data by the demanding type induces convergence to a maxmin best reply to the absolutely continuous models with respect to the true one. In formal terms

$$a \in \operatorname{argmax}_{a' \in A} \min_{p: \exists q \gg p} \mathbb{E}_{p_{a'}} [u(a', y)].$$

In contrast, a statistically sophisticated type maintains a non-trivial concern for misspecification. If their behavior converges, it converges to a robust control best reply to the models closest to the actual data-generating process. Moreover, the misspecification concern is endogenously determined by how well the best models fit the evidence generated by the limit action. In formal terms,

$$a \in \operatorname{argmax}_{a' \in A} \int_Q \min_{p_{a'} \in \Delta(Y)} \left(\mathbb{E}_{p_{a'}} [u(a', y)] + \frac{1}{\lambda} R(p_{a'} || q_{a'}) \right) d\nu(q)$$

with

$$\operatorname{supp} \nu \subseteq \operatorname{argmin}_{q \in Q} R(p_a^* || q_a)$$

and

$$\lambda \simeq \operatorname{argmin}_{q \in Q} R(p_a^* || q_a).$$

Finally, I point out that an endogenous concern for misspecification induces cycles, showing how my model could be used to rationalize the cyclical monetary policy documented in [Sargent 2008].

4. FINAL REMARKS

The comments in Section 2 may lead to the question of whether the elaborate decision criteria proposed by the literature on Decision Theory under uncertainty

are needed, given the theoretical advance in Computer Science. If we can achieve, say, no regret, at an optimal speed under some algorithm, why bother, at least in dynamic environments, with normative Decision Theory?

In my view, there are two main reasons why this perspective is fallacious, both related to the speed of convergence. The environment faced by the economic agents will often change before the positive results of the Computer Science literature have time to apply. Similarly, the agents are often not completely impatient. In all these situations, using models that allow the agent to extrapolate from the consequences of one course of action to another (i.e., that do not treat the environment as a generalized bandit problem) is extremely valuable. At the same time, the complexity of the environment implies those models or scenarios will probably be incomplete and calls for combining them with robustness or cautious evaluations.

REFERENCES

- BANERJEE, S., GKATZELIS, V., GOROKH, A., AND JIN, B. 2022. Online nash social welfare maximization with predictions. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 1–19.
- BATTIGALLI, P., FRANCETICH, A., LANZANI, G., AND MARINACCI, M. 2019. Learning and self-confirming long-run biases. *Journal of Economic Theory* 183, 740–785.
- BREIMAN, L., LECAM, L., AND SCHWARTZ, L. 1964. Consistent estimates and zero-one sets. *The Annals of Mathematical Statistics* 35, 1, 157–161.
- CAMARA, M. K. 2022. Computationally tractable choice. In *EC*. 28.
- CERREIA-VIOGLIO, S., DILLENBERGER, D., AND ORTOLEVA, P. 2015. Cautious expected utility and the certainty effect. *Econometrica* 83, 2, 693–728.
- CERREIA-VIOGLIO, S., HANSEN, L. P., MACCHERONI, F., AND MARINACCI, M. 2022. Making decisions under model misspecification.
- CERREIA-VIOGLIO, S., MACCHERONI, F., MARINACCI, M., AND MONTRUCCHIO, L. 2011. Uncertainty averse preferences. *Journal of Economic Theory* 146, 4, 1275–1330.
- CERREIA-VIOGLIO, S., MACCHERONI, F., MARINACCI, M., AND MONTRUCCHIO, L. 2013. Classical subjective expected utility. *Proceedings of the National Academy of Sciences* 110, 17, 6754–6759.
- CESA-BIANCHI, N. AND LUGOSI, G. 2006. *Prediction, learning, and games*. Cambridge university press.
- DIACONIS, P. AND FREEDMAN, D. 1986. On the consistency of bayes estimates. *The Annals of Statistics* 14, 1, 1–26.
- DOOB, J. L. 1949. Application of the theory of martingales. *Le calcul des probabilités et ses applications*, 23–27.
- ECHENIQUE, F., GOLOVIN, D., AND WIERMAN, A. 2011. A revealed preference approach to computational complexity in economics. In *Proceedings of the 12th ACM conference on Electronic commerce*. 101–110.
- EPSTEIN, L. G. AND SCHNEIDER, M. 2007. Learning under ambiguity. *The Review of Economic Studies* 74, 4, 1275–1303.
- ESPONDA, I. AND POUZO, D. 2016. Berk–nash equilibrium: A framework for modeling agents with misspecified models. *Econometrica* 84, 3, 1093–1130.
- FUDENBERG, D., KLEINBERG, J., LIANG, A., AND MULLAINATHAN, S. 2022. Measuring the completeness of economic models. *Journal of Political Economy* 130, 4, 956–990.
- GHIRARDATO, P., MACCHERONI, F., AND MARINACCI, M. 2004. Differentiating ambiguity and ambiguity attitude. *Journal of Economic Theory* 118, 2, 133–173.
- GILBOA, I. 2009. *Theory of decision under uncertainty*. Number 45. Cambridge university press.
- GILBOA, I. 2023. Decision under uncertainty state of the science. *Annual Review of Economics* forthcoming.

- GILBOA, I. AND MARINACCI, M. 2016. Ambiguity and the bayesian paradigm. In *Readings in formal epistemology: Sourcebook*. Springer, 385–439.
- GILBOA, I. AND SCHMEIDLER, D. 1989. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics* 18, 2, 141–153.
- GUL, F. AND PESENDORFER, W. 2014. Expected uncertain utility theory. *Econometrica* 82, 1, 1–39.
- HANSEN, L. P. AND MARINACCI, M. 2016. Ambiguity aversion and model misspecification: An economic perspective. *Statistical Science* 31, 4, 511–515.
- HANSEN, L. P. AND SARGENT, T. J. 2001. Robust control and model uncertainty. *American Economic Review* 91, 2, 60–66.
- HANSEN, L. P. AND SARGENT, T. J. 2011. Robustness. In *Robustness*. Princeton university press.
- ILUT, C. L. AND SCHNEIDER, M. 2022. Modeling uncertainty as ambiguity: A review.
- LANZANI, G. 2022. Correlation made simple: Applications to salience and regret theory. *The Quarterly Journal of Economics* 137, 2, 959–987.
- LANZANI, G. 2023. Dynamic concern for misspecification. In *Proceedings of the 24th ACM Conference on Economics and Computation*.
- LANZANI, G. 2024. Dynamic concern for misspecification.
- MACCHERONI, F., MARINACCI, M., AND RUSTICHINI, A. 2006. Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica* 74, 6, 1447–1498.
- MAHDIAN, M., NAZERZADEH, H., AND SABERI, A. 2012. Online optimization with uncertain information. *ACM Transactions on Algorithms (TALG)* 8, 1, 1–29.
- MARINACCI, M. 2002. Learning from ambiguous urns. *Statistical Papers* 43, 1, 143–151.
- NOAM NISAN, TIM ROUGHGARDEN, E. T. V. V. V. 2011. *Algorithmic Game Theory*. Cambridge University Press.
- SARGENT, T. J. 2008. Evolution and intelligent design. *American Economic Review* 98, 1, 5–37.
- SAVAGE, L. J. 1954. *The foundations of statistics*. Courier Corporation.
- SCHMEIDLER, D. 1989. Subjective probability and expected utility without additivity. *Econometrica: Journal of the Econometric Society*, 571–587.
- STRZALECKI, T. 2011. Axiomatic foundations of multiplier preferences. *Econometrica* 79, 1, 47–73.
- VALIANT, L. G. 1984. A theory of the learnable. *Communications of the ACM* 27, 11, 1134–1142.

Welfare-Maximizing Pooled Testing

SIMON FINSTER

CREST-ENSAE

and

MICHELLE GONZÁLEZ AMADOR

UNU-MERIT & Maastricht University

and

EDWIN LOCK

University of Oxford

and

FRANCISCO MARMOLEJO-COSSÍO

Harvard University

and

EVI MICHA

University of Toronto

and

ARIEL D. PROCACCIA

Harvard University

This letter provides an overview of our recent work on COVID-19 testing mechanisms that appeared at EC'23. Large-scale testing is crucial in pandemics but resources are often prohibitively constrained. We study a scenario in which a population under lockdown utilizes a limited budget of tests to allow healthy individuals to resume in-person activities. Our work explores the optimal allocation of pooled tests in populations that are heterogeneous with respect to individual infection probabilities and utilities that materialize if included in a negative test (and being permitted to resume in-person activities). Non-overlapping allocations of tests, where no individual in the population is included in more than one pooled test, are both conceptually and logistically simpler to implement. We show that the welfare gain from overlapping testing over non-overlapping testing is bounded. Moreover, we design a heuristic mechanism for finding test allocations that is fast and empirically near-optimal. We also implement our mechanism in practice and provide experimental evidence on the benefits of utility-weighted pooled testing in a real-world setting. Our randomized trial at a higher education research institute in Mexico suggests that performance and mental health outcomes of participants under our testing mechanism are no worse than under the counterfactual of full access for individuals without testing.

Categories and Subject Descriptors: []:

General Terms: Algorithms; Theory; Economics; Experimentation; Performance

Additional Key Words and Phrases: pooled testing, welfare maximization, approximation guarantees, COVID-19, experiment, algorithms

Authors' addresses: simon.finster@ensae.fr, mgonzalez@merit.unu.edu, edwin.lock@cs.ox.ac.uk, fjmarmol@seas.harvard.edu, emicha@cs.toronto.edu, arielpro@seas.harvard.edu

1. INTRODUCTION

A challenging reality of pandemic response is that policymakers are often forced to make decisions with imperfect information. During the COVID-19 pandemic, governments across the globe imposed lockdowns to curb viral spread. These policies resulted from a lack of fine-grained information regarding infection prevalence among the population but came at a high economic and social cost [Deb et al. 2022; Camera and Giofr  2021]. Population testing programs emerged as a viable alternative to blanket lockdowns, allowing individuals who are healthy to resume normal activities. In practice however, finances and resources for testing can be prohibitively constrained. As a result, comprehensive individual testing is often infeasible particularly in low- and middle-income countries [Kavanagh et al. 2020; Dhabaan et al. 2020; Abera et al. 2020].

In this letter, we highlight the main technical contributions and experimental results from our work [Finster et al. 2023]. We also describe insights that helped us bridge theory and practice. Our work focuses on optimally using a limited budget of tests to alleviate the costs of lockdown for a given population under the assumption that individuals who are verifiably healthy are allowed to return to in-person activities. To extend the reach of a limited testing budget, we make use of pooled testing. In a pooled test, samples of multiple individuals are pooled together and tested as one. If this test is positive, we know that at least one individual in the pool is infected; otherwise, all pooled individuals are healthy.

Pooled testing traces back to [Dorfman 1943], who devised the technique in 1943 to screen large numbers of soldiers for syphilis, and has since been the subject of a vast literature.¹ Our point of departure from prior work is the observation that individuals in the given population may have different utilities for resuming in-person activities, as well as different degrees of viral exposure resulting in different probabilities of being infected. The core challenge lies in determining the most effective allocation of tests to maximize the expected utility of individuals capable of resuming in-person activities.

Most importantly, our problem setting and computational techniques are heavily motivated by a collaboration with the Potosinian Institute of Scientific and Technological Research (IPICYT) during the COVID-19 pandemic. IPICYT is a public research institution in the state of San Luis Potos , Mexico. During the latter stages of the pandemic, IPICYT was subjected to a full lockdown. Moreover, researchers and students had differing priorities for returning to on-site facilities for their work. IPICYT had limited resources for qPCR tests but their in-house testing facilities were able to process pooled tests (of up to 5 individuals in each pool). This motivated our research collaboration with IPICYT to develop a pooled testing mechanism that tests those who most urgently need to return to in-person activities.

When given a fixed budget of tests and a population to test, the number of ways to allocate pooled tests is vast. An important practical constraint from our partners at IPICYT was that no individual be included in more than one test in

¹See [Aldridge et al. 2019] and [Du et al. 2000] for extensive surveys on computational techniques in pooled testing.

an allocation (we call this a *non-overlapping* testing allocation), as the logistical overhead of splitting samples into multiple tests was deemed to be prohibitive, and potentially prone to human error. For this reason, our paper starts by quantifying the loss incurred by confining allocations to be non-overlapping. Our theoretical analysis establishes that the worst-case ratio between an optimal overlapping and non-overlapping allocation of tests is bounded by a constant factor, supporting the prioritization of non-overlapping testing. Recognizing the computational complexity of determining an optimal testing allocation (that may include overlapping tests), we introduce a greedy polynomial-time algorithm for computing non-overlapping test allocations that is both conceptually simple and computationally efficient. As we show, it also guarantees an approximate solution with at least one-fifth of the welfare achieved in the optimal non-overlapping testing regime. Moreover, it performs near-optimally in numerical experiments that use real-world data. The greedy algorithm is used in our randomized trial due to its effectiveness.

Our algorithmic testing framework was empirically validated in a small-scale randomized control trial (RCT) conducted in 2022 at IPICYT in Mexico with a population of 130 individuals. The goal of our RCT was to showcase the feasibility of our testing mechanism in practice, and to understand the impact on performance and mental health outcomes compared to full reopening without testing. The trial protocol, ensuring in-person access only for those with negative qPCR test results, curtailed contagion within the institution. Importantly, the results indicated no adverse effects on participants' productivity and mental health, comparing favorably to a counterfactual scenario of unrestricted access without testing and highlighting the effectiveness of the proposed testing approach in practice.

Related Work. During the COVID-19 pandemic, pooled testing emerged as resource-efficient testing strategy [Sanghani et al. 2021; Mutesa et al. 2021; Nalbantoglu 2020]. Our work relates to recent contributions such as [Lipnowski and Ravid 2021], which study optimal testing allocations with respect to welfare maximization in heterogeneous populations. Our work introduces a more general setup, acknowledging diverse health probabilities and utility variations among individuals. We address the complexities of overlapping and non-overlapping testing scenarios, considering practical algorithms for optimal testing allocations that account for heterogeneous utilities. Moreover, [Ely et al. 2021], [Brault et al. 2021], and [Gollier and Gossner 2020] explore diverse aspects of test allocation, from differential costs and sensitivities to early screening and infection prevalence estimation.

2. MAIN THEORETICAL RESULTS

We consider a heterogeneous population in which every individual is characterized by an independent probability of infection and a utility of returning to in-person activities. We also have a limited budget of available pooled tests. In each test, a limited number of samples of individuals are pooled.² The test result is negative if all the samples are negative, otherwise the test is positive. Due to independence,

²Pool sizes in pooled tests are limited due to biological constraints. Our partners in Mexico have replicated techniques from [Sanghani et al. 2021] to achieve a maximal pool size of 5 with saliva samples.

the probability of a negative result is equal to the product of the probabilities of the individuals included in the test being healthy. A *test allocation* indicates the samples of which individuals are assigned to each available test. The expected utility of an individual under a test allocation is equal to their probability of being included in at least one negative test multiplied by their utility to return to in-person activities. The goal is to maximize the *welfare*, i.e. the sum of expected utilities, earned under a test allocation.

Performance of Non-Overlapping Testing. We are particularly interested in *non-overlapping* test allocations, which include each individual in at most one test. In general, overlapping testing can achieve higher welfare than test allocations that are restricted to not overlap. However, non-overlapping test allocations are often strongly preferred for logistical reasons, as was the case with IPICYT. A natural question is to identify how much welfare may be lost by restricting to non-overlapping tests. If the difference in welfare achievable with overlapping and non-overlapping testing is not too large, even institutions with the logistical capacity to run overlapping tests may choose the latter to reduce costs.

Given a fixed population and testing budget, we define the *overlap welfare ratio* as the ratio of the welfare of an optimal test allocation over the welfare of an optimal non-overlapping test allocation. The *gain of overlapping* given a fixed budget is the maximum welfare ratio across all possible populations and this budget. The main question that we aim to answer is:

How large can the gain of overlapping become?

Surprisingly, we find that this gain is a small constant for any value of the budget.

THEOREM 2.1. *For any budget, the gain of overlapping is at most 4.*

Algorithms for Computing Approximately Optimal Testing Allocations. Given Theorem 2.1, we focus on providing efficient algorithms for computing optimal non-overlapping testing allocations. When there is just one test available, we have two efficient algorithms for computing an approximately optimal testing allocation: a modification of the FPTAS of [Goldberg and Rudolf 2020] that accounts for pool size constraints, and a mixed integer conic optimization problem which we solve with commercial solvers in practice. In both cases, the algorithms efficiently compute allocations which achieve a $(1 - \varepsilon)$ proportion of the optimal single test welfare for $\varepsilon > 0$.

When there are more than one test available, we design a greedy algorithm which repeatedly allocates a single test to untested individuals in the population, using our algorithms for one test. We prove that this greedy algorithm provides a constant factor approximation to the optimal non-overlapping test allocation welfare. In practice, we observed that greedy is near-optimal.

THEOREM 2.2. *For any population and testing budget, the greedy algorithm achieves at least $\frac{1-\varepsilon}{5}$ of the welfare of the optimal non-overlapping test allocation, for any $\varepsilon > 0$.*

3. RANDOMIZED CONTROLLED TRIAL

In order to provide causal evidence of the efficacy of the testing mechanism, we designed a two-group randomized controlled trial (RCT). The data collected during the RCT allowed us to a) validate our theoretical results using real-life parameter distributions (health probabilities and utilities), and b) evaluate population welfare (aggregate utilities of returning to onsite work as measured by productivity, performance, learning, stress, subjective well-being). Our results indicate that our algorithm and testing protocol are, indeed, a viable policy tool for pandemic response.

In September 2022, as campus facilities reopened, we conducted the RCT with a population of 130 individuals. The treatment group followed our algorithmic testing strategy, with access to campus granted only upon receiving a negative test result. The control group was granted permission to return to the institute without testing. (The treatment and control groups were instructed not to interact, in order to prevent contagion.) Testing services were provided by the National Laboratory of Agricultural, Medical, and Environmental Biotechnology (LANBAMA), a testing facility within IPICYT.

We worked with researchers at IPICYT to gather population data to estimate individuals' utilities for in-person access and their health probabilities. In collaboration with epidemiologists at IPICYT and in the local state of San Luis Potosí we were able to achieve reliable estimates of infection probabilities. We also allowed individuals to express onsite work preferences for two-day windows through the allocation of a virtual token budget. This helped us avoid scheduling individuals for testing on days they did not wish to access IPICYT facilities in the first place, and allocate more tests to particular days that proved more popular.

To coordinate the RCT, we developed a [web application](#). The web app included intake and outtake surveys for treatment and control participants, from which we collected information with regards to their social and economic well-being, and their preferences and beliefs about their subjective well-being. These data were used to i) estimate static utilities and ii) evaluate the (static and dynamic) population impact of the testing mechanism. Through the web app we also managed their token budget for weekly onsite work preferences, communicated weekly testing schedules, and allowed the lab to anonymously communicate test results to treatment participants.

Evaluation and methods. In the RCT, we measured subjects' stress levels and subjective well-being (life satisfaction), as well as self-assessed performance, productivity, and learning. We obtain these measures through survey questions that subjects are invited to answer before (baseline) and after (endline) the trial period. The treatment effect on the treated was estimated with bivariate linear regressions, based on static scores as well as first differences³, using the above-mentioned outcomes as dependent variables. The trial period was of low viral prevalence in San Luis Potosí. The general decrease in contagion rates was reflected in our study:

³We also performed equivalence tests, and multinomial logistic regression models for non-normally distributed outcomes. These robustness checks corroborate our results from our preferred model specifications. We also collected a number of covariates for further robustness checks and heterogeneity analyses of our estimations.

only one pool tested positive in the treatment group. No individuals self-reported having experienced symptoms in the control group. This low number means that we are unable to make any strong claims regarding the protocol's health benefits.

We find no statistical evidence that our protocol has a negative effect on participants' work/study performance, learning, or mental health, despite the increased effort in coordination it demands from them compared to a full reopening (the protocol followed by the control group). At the same time, our strategy ensures greater safety for all participating individuals compared to a full reopening without any safety mechanisms in place. We conjecture that accounting for welfare is the crucial ingredient in our mechanism, enabling in-person access for those who need and benefit from it the most.

4. BRIDGING RESEARCH AND PRACTICE

The social relevance of our research lies in its successful implementation and population impact, grounded in strong theoretical foundations. [Lock et al. 2021] highlight that a crucial component of creating successful research-to-practice pipelines is establishing strong partnerships with local regulatory agencies; these partnerships guarantee that applied research projects are in congruence with local realities, both in terms of administrative constraints and knowledge of the local flow of know-how.

In the case of Mexico, our main partners were the state research councils belonging to Mexico's National Network of State Councils of Science and Technology ([REDNACECYT](#))⁴. The support of the local council of San Luis Potosí, [COPOCYT](#), was instrumental in facilitating introductions and encouraging collaborations between local researchers and our team. It provided an important degree of trust when working with researchers at IPICYT and beyond, such as when we reached out to local epidemiologists at the National Autonomous University of San Luis Potosí ([UASLP](#)) to help inform local infection rates to be used in the randomized trial of our protocol.

Moreover, as the pandemic progressed, COPOCYT also played an important role in helping our project remain aligned with the changing nature of local administrative constraints. This was especially important with regards to government policy around testing, as our protocol made heavy use of novel pooled testing lab methodology, for which health and regulatory policies were still being developed; their support also extended to the alignment of our project to changing national and local policies regarding lockdown regulations, which continued to evolve alongside the pandemic. Ultimately, our communication with these local partners helped us maintain a positive feedback loop. This ensured that the implementation of our protocol was as effective as possible while avoiding unnecessary complexity.

5. FUTURE WORK

Our pooled testing framework belongs to a growing algorithmic literature on resource allocation that incorporates population heterogeneity. Our testing protocol

⁴REDNACECYT and its constituent councils operate under the purview of the Mexican National Council of Science and Technology ([CONACYT](#)), which in turn serves as Mexico's primary government entity in charge of the promotion, funding and regulation of scientific and technological activities.

offers institutions an economical and secure solution to safeguard their entire community.

Looking ahead, there are several immediate avenues for future exploration. On the theoretical front, there is a gap between our upper bound of 4 and lower bound of $7/6$ on the overlap welfare ratio, as well as an upper bound of 5 on the approximation factor of the greedy algorithm. We anticipate that even tighter bounds are attainable when utilities are confined to a fixed number of values, such as in dichotomous or trichotomous populations.

The testing and re-integration policy we advocate is static, assuming a one-shot setting where a testing budget is fully utilized by a policymaker. Dynamic testing, on the other hand, could adjust allocations adaptively based on previous test results. Furthermore, policymakers may have access to different types of tests, each with varying costs and performance metrics (e.g., pool size and sensitivity). Determining optimal budget-constrained allocations in this heterogeneous test setting remains an open question.

Crucially, we hope that the insights gained from the performance and efficacy of our welfare-maximizing testing mechanism will aid in better protecting resource-constrained communities during infectious disease outbreaks. Our pooled testing protocol may extend to mass screening for HIV/AIDS and prove equally significant in pooled frameworks for organ donation.

REFERENCES

- ABERA, A., BELAY, H., ZEWUDE, A., GIDEY, B., NEGA, D., DUFERA, B., ABEBE, A., ENDRIYAS, T., GETACHEW, B., BIRHANU, H., ET AL. 2020. Establishment of COVID-19 testing laboratory in resource-limited settings: challenges and prospects reported from Ethiopia. *Global Health Action* 13, 1, 1–7.
- ALDRIDGE, M., JOHNSON, O., AND SCARLETT, J. 2019. Group Testing: An Information Theory Perspective. *Foundations and Trends® in Communications and Information Theory* 15, 3-4, 196–392.
- BRAULT, V., MALLEIN, B., AND RUPPRECHT, J.-F. 2021. Group testing as a strategy for COVID-19 epidemiological monitoring and community surveillance. *PLoS computational biology* 17, 3, 1–25.
- CAMERA, G. AND GIOFFRÉ, A. 2021. The economic impact of lockdowns: A theoretical assessment. *Journal of Mathematical Economics* 97, 1–14.
- DEB, P., FURCERI, D., OSTRY, J. D., AND TAWK, N. 2022. The economic effects of covid-19 containment measures. *Open Economies Review* 33, 1, 1–32.
- DHABAAN, G. N., AL-SONEIDAR, W. A., AND AL-HEBSHI, N. N. 2020. Challenges to testing COVID-19 in conflict zones: Yemen as an example. *Journal of global health* 10, 1, 1–3.
- DORFMAN, R. 1943. The detection of defective members of large populations. *The Annals of Mathematical Statistics* 14, 4, 436–440.
- DU, D., HWANG, F. K., AND HWANG, F. 2000. *Combinatorial group testing and its applications*. Vol. 12. World Scientific, Singapore.
- ELY, J., GALEOTTI, A., JANN, O., AND STEINER, J. 2021. Optimal test allocation. *Journal of Economic Theory* 193, 1–17.
- FINSTER, S., AMADOR, M. G., LOCK, E., COSSIO, F. M., MICHA, E., AND PROCACCIA, A. 2023. Welfare-maximizing pooled testing. In *ACM Conference on Economics and Computation 2023*. EC '23. Association for Computing Machinery, New York, NY, USA.
- GOLDBERG, N. AND RUDOLF, G. 2020. On the complexity and approximation of the maximum expected value all-or-nothing subset. *Discrete Applied Mathematics* 283, 1–10.
- GOLLIER, C. AND GOSSNER, O. 2020. Group testing against Covid-19. Tech. rep., EconPol Policy Brief.

- KAVANAGH, M. M., ERONDU, N. A., TOMORI, O., DZAU, V. J., OKIRO, E. A., MALECHE, A., ANIEBO, I. C., RUGEGE, U., HOLMES, C. B., AND GOSTIN, L. O. 2020. Access to lifesaving medical resources for African countries: COVID-19 testing and response, ethics, and politics. *The Lancet* 395, 10238, 1735–1738.
- LIPNOWSKI, E. AND RAVID, D. 2021. Pooled testing for quarantine decisions. *Journal of Economic Theory* 198, 1–30.
- LOCK, E., MARMOLEJO-COSSÍO, F. J., JONNERBY, J., RAJGOPAL, N., GUZMÁN-GUTIÉRREZ, H. A., BENAVIDES-VÁZQUEZ, L. A., TELLO-AYALA, J. R., AND LAZOS, P. 2021. Optimal testing and containment strategies for universities in mexico amid covid-19. In *Equity and Access in Algorithms, Mechanisms, and Optimization*. Proceedings of EAAMO’21, USA, 1–9.
- MUTESA, L., NDISHIMYE, P., BUTERA, Y., SOUOPGUI, J., UWINEZA, A., RUTAYISIRE, R., NDORICIMPAYE, E. L., MUSONI, E., RUJENI, N., NYATANYI, T., ET AL. 2021. A pooled testing strategy for identifying SARS-CoV-2 at low prevalence. *Nature* 589, 7841, 276–280.
- NALBANTOGLU, O. U. 2020. Group testing performance evaluation for SARS-CoV-2 massive scale screening and testing. *BMC Medical Research Methodology* 20, 1, 1–11.
- SANGHANI, H. R., NAWROT, D. A., MARMOLEJO-COSSÍO, F., TAYLOR, J. M., CRAFT, J., KALIMERIS, E., ANDERSSON, M. I., AND VASUDEVAN, S. R. 2021. Concentrating Pooled COVID-19 Patient Lysates to Improve Reverse Transcription Quantitative PCR Sensitivity and Efficiency. *Clinical Chemistry* 67, 5, 797–798.

Leveraging Reviews: Learning to Price with Buyer and Seller Uncertainty

WENSHUO GUO and NIKA HAGHTALAB

University of California, Berkeley

and

KIRTHEVASAN KANDASAMY

University of Wisconsin, Madison

and

ELLEN VITERCIK

Stanford University

Customers can access hundreds of reviews for a single product in online marketplaces. Buyers often use reviews from other customers that share their type—such as height for clothing or skin type for skincare products—to estimate their values, which they may not know a priori. Customers with few relevant reviews may hesitate to purchase except at a low price, so for the seller, there is a tension between setting high prices and ensuring that there are enough reviews so buyers can confidently estimate their values. Simultaneously, sellers may use reviews to gauge the demand for items they wish to sell. In this work, we study this pricing problem in an online setting where the seller interacts with a set of buyers of finitely many types, one by one, over a series of T rounds. At each round, the seller first sets a price. Then, a buyer arrives and examines the reviews of the previous buyers with the same type, which reveal those buyers' ex-post values. Based on the reviews, the buyer decides to purchase if they have good reason to believe their ex-ante utility is positive. Crucially, the seller does not know the buyer's type when setting the price, nor even the distribution over types. We provide a no-regret algorithm that the seller can use to obtain high revenue. When there are d types, after T rounds, our algorithm achieves a problem-independent $\tilde{O}(T^{2/3}d^{1/3})$ regret bound. However, when the smallest probability $q_{\min}$ that any given type appears is large, specifically when $q_{\min} \in \Omega(d^{-2/3}T^{-1/3})$, the same algorithm achieves a $\tilde{O}(T^{1/2}q_{\min}^{-1/2})$ regret bound. We complement these upper bounds with matching lower bounds in both regimes, showing that our algorithm is minimax optimal up to lower-order terms.

This is a summary of work that won the *Exemplary AI Track Paper Award* at EC'24.

Categories and Subject Descriptors: 500 [Theory of computation]: Computational pricing and auctions; 500 [Theory of computation]: Online learning theory

General Terms: Algorithms; Economics; Theory

Additional Key Words and Phrases: Pricing, Reviews, Revenue maximization, Online learning

1. INTRODUCTION

The rapid growth of e-commerce has allowed customers to gain insights from thousands of reviews before deciding whether to purchase an item. Customers often use reviews by buyers who share their “type” —such as body type for clothes or skin type for skincare products—to develop high-fidelity estimates of their values

Authors' addresses: {wsguo, nika}@berkeley.edu, nika@berkeley.edu, kandasamy@cs.wisc.edu, vitercik@stanford.edu.

for items, which are quantities they may be uncertain of before purchasing.

When learning from reviews, a customer’s purchase decision is not just a function of the item’s price but also of how certain the customer is about her valuation, which in turn depends on the earlier sales and reviews of the items. This leads to a tension between setting revenue-optimal prices while ensuring buyers have enough reviews to estimate their values confidently. This tension is perhaps most clear for customers of rare types (for example, particularly tall or short individuals shopping for clothing) who may find only a few reviews from similar customers and, due to risk aversion, may only be willing to buy at relatively low prices.

We introduce a model that simultaneously captures the seller’s pricing problem, the buyers’ learning problem, and the modus through which the buyers learn: reviews. We study how a seller—uncertain about the buyers’ type distribution—can learn to set high-revenue prices when the buyers themselves are unsure about their values and are learning from reviews. Thus, there is information uncertainty on both sides of the market: the seller is uncertain about which buyer will arrive and the buyers’ type distribution, but the buyer, who knows their type, is unsure about their *ex-ante* value. Both sides of the market are operating with significantly less information than has historically been assumed in mechanism design.

We study this pricing problem with an online sequential learning model where the seller attempts to sell identical copies of an item to a series of distinct buyers over T timesteps. Each buyer has one of d types drawn from a distribution $\mathcal{P}$, and a buyer of type i has an *ex-ante* value of θ_i for the item. At each timestep t , the seller sets a price p_t . The seller knows the *ex-ante* values $\theta_1, \dots, \theta_d$ and thus has some limited information about the buyers (for example, from market research), but he does not know the buyer’s type on each round nor even the distribution $\mathcal{P}$. If a buyer of type i purchases the item, they will leave a review communicating their *ex-post* value for the item, which is a random variable with mean θ_i . To decide whether to purchase, a new buyer evaluates reviews left by buyers of type i who bought the item in the past. Specifically, the buyer at round $t \in [T]$ uses the past reviews to select a threshold τ_t and chooses to buy as long as $p_t \leq \tau_t$. If the buyer’s threshold τ_t is too pessimistic—for example, it always equals zero no matter the reviews—then optimizing revenue would be hopeless. In our model, we bound the level of risk-aversion that the buyer can display: we assume that τ_t is at least a lower confidence bound we denote LB_t that equals the average of the reviews left by buyers with the same type, minus an uncertainty term that depends on the number of such reviews. Intuitively, the buyer can be confident that their *ex-ante* value is at least LB_t with high probability, so they always buy if they have good reason to believe that their *ex-ante* utility will be positive.

The *ex-post* value is the actual experience of the buyer and is different from the *ex-ante* value due to exogenous stochastic factors that cannot be known at the time of purchase (for example, manufacturing defects, color on the website not matching the actual color). Hence, when there is complete information, the buyer decides based on their *ex-ante* value. In our problem, the buyer does not even know their *ex-ante* value and uses reviews from previous buyers to estimate it.

We provide a no-regret learning algorithm for the seller that balances setting high-revenue prices with soliciting reviews from rare but high-value customers. The key

challenge is that the seller does not know the current buyer's type on each round *a priori*: the prices are anonymous. This means the seller does not know the number of reviews that the buyer will use to construct their value estimate. A buyer on any round could be (i) a high-value type, but uncertain of their value since their type has few reviews, and thus may be hesitant to purchase except at a low price, (ii) a high-value type, and more confident of their value since their type has many reviews, and thus is willing to purchase at a high price, or (iii) a low-value type whom the seller should not target even if they were absolutely sure of their value since it leads to small per-purchase revenue. In the first case, it may be worthwhile to initially set a low price to solicit enough reviews to ensure future purchases at a higher price, winning over these rare but high-value customers. The seller, however, has to decide which buyers to win over without knowing the buyer's type on each round, nor even the distribution over types. He may, therefore, wastefully offer a low price to a buyer who would be willing to buy at a higher price.

2. RELATED WORK

Learning to price when buyers do not know their values. Several papers have studied selling repeatedly to a single buyer while the buyer is learning from their experience [Papadimitriou et al. 2022; Ashlagi et al. 2016; Chawla et al. 2022; Feng et al. 2018; Weed et al. 2016; Kandasamy et al. 2023]. However, buyers on online platforms often do not return repeatedly to buy the same item and can only obtain feedback from previous buyers via reviews. In this paper, we study a setting where the buyers can only learn from past reviews.

Ifnach et al. [2019] consider a similar pricing problem for the seller when the buyers learn from reviews. However, their model is limited to one buyer type, whereas we study the setting with multiple buyer types. Moreover, the seller does not know the frequency of each type and the type of buyer who arrives at each round, which leads to crucial difficulties in our analysis.

Learning to price when buyers know their values. Zhao and Chen [2020] study a setting where the buyers know their values, but the seller does not know the distribution over buyers' values. They present an algorithm that uses buyer reviews to obtain a $\tilde{O}(T^{1/2})$ regret bound. In contrast, if the seller only observes purchase decisions and not reviews, Kleinberg and Leighton [2003] provide a $\Omega(T^{2/3})$ lower bound. While they show that this bound can be improved to $\tilde{\Theta}(T^{1/2})$, it requires additional distributional assumptions.

Selling to no-regret buyers who know their values. When buyers know their values, the buyer may strategically improve their purchase decisions or bidding strategy over repeated interactions to achieve a higher accumulated utility. No-regret learning has been explored as a model of buyer behavior [Braverman et al. 2018; Deng et al. 2019; Nekipelov et al. 2015; Devanur et al. 2014]. In this literature, buyers know their values, whereas we work with buyers who do not and need to estimate their values from historical reviews. This leads to different dynamics.

Buyers' social learning from reviews. Our work is related to a rich literature on buyer behavior and social learning from reviews when buyers do not know their values [Ifnach et al. 2019; Boursier et al. 2022; Han and Anderson 2020; Chamley

2004; Besbes and Scarsini 2018; Bose et al. 2006; Crapis et al. 2017; Kakhbod et al. 2021; Acemoglu et al. 2022]. Much of this research can be categorized into two groups depending on whether the decision model is Bayesian or non-Bayesian. It may be computationally challenging for buyers to compute Bayesian updates, so several papers relax this assumption [Crapis et al. 2017; Besbes and Scarsini 2018]. Besbes and Scarsini [2018], for example, study both Bayesian buyers and buyers with limited rationality who can only observe the average of the past reviews. They analyze the conditions under which buyers can recover a product’s true quality based on observed feedback. Unlike our paper, the buyers have private signals about the item, influencing their purchase decisions. Our model can be seen as situated between these two extremes because the purchase decisions depend on the average of the past reviews and the number of those reviews. Moreover, whereas Besbes and Scarsini [2018] analyze risk-neutral buyers, we study a form of risk aversion where buyers may not purchase even if the price is below the average reviews.

Unlike this prior research, we do not assume all buyers share a specific decision policy: we identify a broad family of decision policies under which our results hold. We only require that the buyer purchases the item if the price is sufficiently low.

3. NOTATION AND ONLINE LEARNING SETUP

In our model, an item is sold repeatedly to a sequence of distinct buyers over T rounds. Each buyer has a type $i \in [d]$, and there is an unknown distribution $\mathcal{P}$ over the types $[d]$. We use the notation $q_i = \Pr_{j \sim \mathcal{P}}[j = i]$ and $q_{\min} = \min_{i \in [d]} q_i$.

The *ex-ante* value of a buyer with type $i \in [d]$ is $\theta_i \in [0, 1]$ and their *ex-post* value is drawn from a distribution $\mathcal{D}_i$ with support $[0, 1]$ and mean θ_i , with $\theta_1 \leq \theta_2 \leq \dots \leq \theta_d$. The seller knows $\theta_1, \dots, \theta_d$ but not $\mathcal{P}, \mathcal{D}_1, \dots, \mathcal{D}_d$. At each round $t \in [T]$:

- (1) There is a set σ_{t-1} of reviews describing past buyers’ types and *ex-post* values.
- (2) The seller first sets a price $p_t \in [0, 1]$.
- (3) A buyer arrives with type $i_t \sim \mathcal{P}$. They observe past reviews of buyers with type i_t , i.e. $\Phi_{i_t, t} = \{v : (i, v) \in \sigma_{t-1} \text{ and } i = i_t\}$, and decide whether to purchase the item. We describe the buyer’s purchasing model in more detail later in this section. The seller is unaware of the buyer’s type i_t when they set the price.
- (4) If the buyer purchases the item, they pay p_t and leave a review of (i_t, v_t) describing both their type and their *ex-post* value $v_t \sim \mathcal{D}_{i_t}$. In this case, $\sigma_t = \sigma_{t-1} \cup \{(i_t, v_t)\}$, and otherwise, $\sigma_t = \sigma_{t-1}$.

Our assumptions and model reflect practical e-commerce settings. First, quite often, it is reasonable to assume that sellers know customers’ *ex-ante* values as they may have inside information. For instance, a skincare product vendor may know that a particular product works better on some skin types. However, buyers may not simply trust the seller if they were to publish this value, as the seller has every incentive to overstate this value to maximize revenue. A buyer would instead decide if a product is suitable for her via independent reviews from other customers. Second, for fairness reasons, in e-commerce platforms, sellers typically have to publish a single price for all customers and cannot sell the item at individualized prices. Third, if a buyer does not purchase an item, they will not leave a review, and the seller has no way of knowing their type or *ex-post* value.

Buyers' purchasing model. At time step t , the agent's purchase decision is defined by a threshold $\tau_t(\sigma_{t-1}, i_t) \geq 0$ that takes as input their type i_t and the reviews left by past agents. Intuitively, $\tau_t(\sigma_{t-1}, i_t)$ represents the agent's estimate of their value θ_{i_t} based on past reviews. The agent purchases the item if $p_t \leq \tau_t(\sigma_{t-1}, i_t)$.

A conservative agent would choose $\tau_t(\sigma_{t-1}, i_t)$ to be low in order to always guarantee that $\tau_t(\sigma_{t-1}, i_t) \leq \theta_{i_t}$, so that they only purchase when their *ex-ante* utility is non-negative. An extreme example of this type of conservatism would set $\tau_t(\sigma_{t-1}, i_t) = 0$, meaning that the agent would only purchase the item if offered for free. Optimizing revenue with such a conservative agent would be hopeless. Therefore, we impose the following natural lower bound on $\tau_t(\sigma_{t-1}, i_t)$:

Definition 3.1. Let Φ_t be the reviews left by agents with type i_t , i.e., $\Phi_t = \{v : (i, v) \in \sigma_{t-1}, i = i_t\}$. Let LB_t be the average minus a standard confidence term:

$$\text{LB}_t = \begin{cases} 0 & \text{if } \Phi_t = \emptyset, \\ \max \left\{ 0, \frac{1}{|\Phi_t|} \sum_{v \in \Phi_t} v - \sqrt{\frac{1}{2|\Phi_t|} \ln \frac{t}{\eta}} \right\} & \text{else.} \end{cases}$$

We say that the agent on round t is η -pessimistic if, $\tau_t(\sigma_{t-1}, i_t) \geq \text{LB}_t$.

This uncertainty term corresponds to the standard Hoeffding confidence interval. Intuitively, as a buyer sees more reviews from his type, he is more certain about his *ex-ante* value. The $\ln t$ term is necessary to construct a valid confidence interval for an arbitrary algorithm as the data may not be independent: the algorithm's price may depend on previous reviews, which will affect future buyers and reviews. This $\ln t$ term is not fundamental—the lower bound does not use it.

Intuitively, the agents can be confident that *regardless* of the policy used by the seller, with probability $1 - \eta$, for all rounds $t \in [T]$, $\theta_{i_t} \geq \text{LB}_t$. Therefore, if the price is lower than LB_t , an η -pessimistic agent will buy the item as they can be confident, based on past reviews, that their *ex-ante* utility $\theta_{i_t} - p_t$ will be non-negative. This restriction bounds the level of pessimism that the agents can display and thus makes it possible to set reasonable prices.

4. ALGORITHM OVERVIEW AND ANALYSIS

This section describes our algorithm, which has two phases. In the first phase, the algorithm sets a price of 0 for $t_0 = \tilde{\Theta}(T^{1/3}d^{2/3})$ rounds. The agent will buy the item at each round since the price is 0 and leave a review. This allows the algorithm to obtain i.i.d. samples from the type distribution $\mathcal{P}$. In phase 2 (i.e., the remaining $T - t_0$ rounds), the algorithm will ignore types that appeared too rarely during phase 1. Intuitively, customers of these types have a low probability of appearance and thus will have more uncertainty about their values due to fewer reviews. The uncertainty term will cause the lower confidence bound LB_t in Definition 3.1 to be small. As the seller will have to choose a low price to target these customers (even if their *ex-ante* value is large), they may have to forego higher revenue from more frequent customer types. Therefore, it is not worthwhile for the algorithm to target these customers. We use Q to denote the buyer types that appeared on a sufficiently large fraction of rounds, as in Figure 1a.

To describe the second phase, we use $\text{rev}(p, Q) = p \Pr_{i \sim \mathcal{P}} [\theta_i \geq p \text{ and } i \in Q]$ to denote the expected revenue of a price p restricted to buyers in Q and $p^*(Q) =$




Fig. 1: Illustration of our algorithm with three types: a red circle, green square, and orange star.

$\arg\max \text{rev}(p, Q)$. In this phase, our algorithm ignores the rare buyers not in Q and aims to set prices that compete with $p^*(Q)$. By competing with $p^*(Q)$, we show that our algorithm also competes with the optimal price p^* .

Observe that $p^*(Q) = \theta_{i_Q}$ for some $i_Q \in Q$. On each round $t > t_0$ of the second phase, our algorithm maintains a set S_t of "active types" such that i_Q is likely in S_t . The algorithm sets the price p_t low enough to ensure that if the current type i_t is in S_t , then the buyer will buy, as in Figure 1b. In particular, we define LB_{it} as the largest price the seller can set to ensure a purchase from a buyer of type i . We then set the price p_t to be the smallest LB_{it} of any type $i \in S_t$.

Next, for each active type $i \in S_t$, the algorithm estimates $\text{rev}(\theta_i, Q)$ along with upper and lower confidence bounds $\hat{\mu}_{i,t}$ and $\check{\mu}_{i,t}$. The algorithm defines $i_0 = \min \{i \in S_t : \hat{\mu}_{i,t} \geq \max_{k \in S_t} \check{\mu}_{k,t}\}$ to be the smallest active type such that θ_{i_0} may plausibly be $p^*(Q)$: for all $i < i_0$, the upper confidence bound on $\text{rev}(\theta_i, Q)$ is small ($\hat{\mu}_{i,t} < \max_{k \in S_t} \check{\mu}_{k,t}$), so it is unlikely that $\theta_i = p^*(Q)$. The algorithm concludes round t by eliminating all types $i < i_0$ from the active set.

Regret definition. We define regret as the difference between:

- (1) The algorithm's total expected revenue, and
- (2) (*baseline*) The expected revenue of the optimal fixed price if the agents bought whenever their *ex-ante* value was larger than the price.

Under the baseline that we compete with, both the buyer and the seller are equipped with more information than in the learning problem: the seller knows all distributions $\mathcal{P}, \mathcal{D}_1, \dots, \mathcal{D}_d$ and the buyers know their *ex-ante* values $\theta_1, \dots, \theta_d$. Therefore, the seller knows *a priori* which customers to target to maximize revenue. Moreover, since the buyers do not need to learn their *ex-ante* values from reviews, the seller can extract higher revenue than they could from uncertain buyers who may only buy when the price is likely lower than their *ex-ante* value. Formally, let $b_t \in \{0, 1\}$ indicate whether the buyer bought on round t and let $p^* = \arg\max p \Pr_{i \sim \mathcal{P}}[\theta_i \geq p]$ be the price with highest expected revenue if the agents bought whenever their *ex-ante* value was larger than the price. Regret is defined as $\mathbb{E}[R_T] = Tp^* \Pr_{i \sim \mathcal{P}}[\theta_i \geq p^*] - \mathbb{E}[\sum p_t b_t]$.

Regret upper bound and proof overview. We contend with several sources of regret. The first phase of the algorithm, where the item is sold for free, inevitably leads to regret, so it must be made as brief as possible. The algorithm then completely disregards the buyer types that appeared too rarely during that phase. This

results in a subset $Q \subseteq [d]$ of buyer types that appear sufficiently often. In the second phase, the algorithm only attempts to optimize revenue with respect to the buyers in Q instead of the entire set $[d]$, which contributes to regret. Finally, the buyers themselves do not know their *ex-ante* values, whereas, under our regret benchmark, buyers buy whenever their *ex-ante* value is larger than the price.

We obtain our final regret bound by analyzing these three sources of error. Our bound depends on the smallest probability that any given type appears, which we denote as $q_{\min}$. If $q_{\min}$ is not tiny—specifically, $q_{\min} > 2d^{-2/3}T^{-1/3}$ —then we obtain a regret bound that scales with $\sqrt{T}$, namely $\tilde{O}(T^{1/2}q_{\min}^{-1/2} + T^{1/3}d^{2/3})$. Otherwise, for arbitrary $q_{\min}$, our regret bound scales with $T^{2/3}$ as $\tilde{O}(T^{2/3}d^{1/3} + T^{1/3}d^{2/3})$.

Regret lower bound and proof overview. Typical bandit lower bounds rely on hypothesis testing arguments to show that any algorithm would struggle to distinguish between similar problems but with different optimal outcomes. Such an analysis would not capture the main difficulty in our setting: how fast customers estimate their *ex-ante* values from past reviews. Instead, our proof leverages the buyers' uncertainty to establish a $\tilde{\Omega}(T^{2/3}d^{1/3})$ worst-case lower bound and a $\tilde{\Omega}(T^{1/2}q_{\min}^{-1/2})$ lower bound when $q_{\min}$ is large. This proves our regret bound is optimal.

Our proof constructs a hard instance where buyer types with a low probability of appearance have comparable *ex-ante* values to types with a high probability of appearance. In each round, an algorithm should decide whether it will target low-probability customers who may be less certain about their value due to fewer reviews and consequently have small LB_t . Keeping prices low to do so leads to low revenue in the current round, but ignoring low-probability customers by choosing a high price risks losing potentially high future revenue. We obtain a tight lower bound by carefully choosing the probability of appearance in our construction.

5. FUTURE DIRECTIONS

Many questions remain open. For example, we assumed that purchases always come with a noisy review. A challenging next step would be to develop pricing strategies when reviews are left with varying probabilities, mimicking real-world buyers.

We studied myopic buyers who make purchase decisions based on value estimates from historical reviews, regardless of the seller's policy. What if the buyers appear over several rounds and behave strategically to lower future prices?

We take a frequentist perspective on this problem. It is also possible to take a Bayesian view of this problem and impose a prior on the *ex-ante* value so that the buyer starts with some prior information. We expect adapting our main proof intuitions to that setting is possible. The main differences would be: (i) we would use Bayesian credible intervals instead of frequentist confidence intervals for the η -pessimism definition, (ii) we would control the Bayes' risk when estimating the *ex-ante* values instead of frequentist concentration arguments, and (iii) our final regret could have a nuanced dependence on this prior which may offer tighter bounds.

Another direction would be to explore the case where the seller does not know the buyers' *ex-ante* values. The key challenge would be related to the regret benchmark: we compete with the optimal price if each buyer knew their *ex-ante* value and bought whenever it was above the price (thus, the buyers are not learning). To compete with this benchmark, we require unbiased estimates of the revenue of

different prices if the buyers bought when their *ex-ante* value was above the price. Computing these unbiased estimates is challenging: if a buyer does not buy on a given round, the algorithm does not learn their type, so it cannot tell whether the buyer has a low *ex-ante* value or he has a high value but a low confidence bound. We can circumvent this subtle challenge if the seller knows the buyers' *ex-ante* values. However, this is not possible if the *ex-ante* values are unknown.

REFERENCES

- ACEMOGLU, D., MAKHDOUNI, A., MALEKIAN, A., AND OZDAGLAR, A. 2022. Learning from reviews: The selection effect and the speed of learning. *Econometrica* 90, 6, 2857–2899.
- ASHLAGI, I., DASKALAKIS, C., AND HAGHPANAH, N. 2016. Sequential mechanisms with ex-post participation guarantees. In *ACM Conference on Economics and Computation (EC)*.
- BESBES, O. AND SCARSINI, M. 2018. On information distortions in online ratings. *Operations Research* 66, 3, 597–610.
- BOSE, S., OROSEL, G., OTTAVIANI, M., AND VESTERLUND, L. 2006. Dynamic monopoly pricing and herding. *The RAND Journal of Economics* 37, 4, 910–928.
- BOURSIER, E., PERCHET, V., AND SCARSINI, M. 2022. Social learning in non-stationary environments. In *International Conference on Algorithmic Learning Theory (ALT)*. 128–129.
- BRAVERMAN, M., MAO, J., SCHNEIDER, J., AND WEINBERG, M. 2018. Selling to a no-regret buyer. In *ACM Conference on Economics and Computation (EC)*.
- CHAMLEY, C. 2004. *Rational herds: Economic models of social learning*. Cambridge University Press.
- CHAWLA, S., DEVANUR, N. R., KARLIN, A. R., AND SIVAN, B. 2022. Simple pricing schemes for consumers with evolving values. *Games and Economic Behavior* 134, 344–360.
- CRAPIS, D., IFRACH, B., MAGLARAS, C., AND SCARSINI, M. 2017. Monopoly pricing in the presence of social learning. *Management Science* 63, 11, 3586–3608.
- DENG, Y., SCHNEIDER, J., AND SIVAN, B. 2019. Prior-free dynamic auctions with low regret buyers. In *Conference on Neural Information Processing Systems (NeurIPS)*.
- DEVANUR, N. R., PERES, Y., AND SIVAN, B. 2014. Perfect Bayesian equilibria in repeated sales. In *Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.
- FENG, Z., PODIMATA, C., AND SYRGKANIS, V. 2018. Learning to bid without knowing your value. In *ACM Conference on Economics and Computation (EC)*.
- HAN, S. AND ANDERSON, C. K. 2020. Customer motivation and response bias in online reviews. *Cornell Hospitality Quarterly* 61, 2, 142–153.
- IFRACH, B., MAGLARAS, C., SCARSINI, M., AND ZSELEVA, A. 2019. Bayesian social learning from consumer reviews. *Operations Research* 67, 5, 1209–1221.
- KAKHBOD, A., LANZANI, G., AND XING, H. 2021. Heterogeneous Learning in Product Markets. *Available at SSRN 3961223*.
- KANDASAMY, K., GONZALEZ, J. E., JORDAN, M. I., AND STOICA, I. 2023. VCG mechanism design with unknown agent values under stochastic bandit feedback. *Journal of Machine Learning Research* 24, 53, 1–45.
- KLEINBERG, R. AND LEIGHTON, T. 2003. The value of knowing a demand curve: Bounds on regret for online posted-price auctions. In *Symposium on Foundations of Computer Science (FOCS)*.
- NEKIPELOV, D., SYRGKANIS, V., AND TARDOS, E. 2015. Econometrics for learning agents. In *ACM Conference on Economics and Computation (EC)*.
- PAPADIMITRIOU, C., PIERRAKOS, G., PSOMAS, A., AND RUBINSTEIN, A. 2022. On the complexity of dynamic mechanism design. *Games and Economic Behavior*.
- WEED, J., PERCHET, V., AND RIGOLLET, P. 2016. Online learning in repeated auctions. In *Conference on Learning Theory (COLT)*.
- ZHAO, H. AND CHEN, W. 2020. Stochastic one-sided full-information bandit. In *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML PKDD)*.

Inequality and Market Design

PIOTR DWORCZAK

Northwestern University & Group for Research in Applied Economics

Policymakers are often concerned about inequalities in the markets they control. In this letter, I argue that mechanism design has not responded sufficiently to the need for a comprehensive theory of inequality-aware market design. I review some of my recent work trying to fill this gap and identify research directions where input from computer scientists would be particularly useful.

Categories and Subject Descriptors: Computer Applications [**Social and Behavioral Sciences**]: Economics

General Terms: Design, Economics, Theory

Additional Key Words and Phrases: Market Design, Mechanism Design, Inequality

1. INTRODUCTION

Within both economics and computer science, mechanism design has been successfully employed as a methodology to understand the optimal design of markets, auctions, trading platforms, and other economic systems. Following the seminal work of Myerson [1981], thousands of papers have been written on revenue-maximizing mechanisms in a variety of contexts and under various constraints. The question of efficient design received a similar level of attention: The classical results about the implementability of efficient outcomes (Vickrey [1961], Clarke [1971], Groves [1973]) have been examined from every possible angle.

Policymakers, though, often attempt to maximize an objective function that evidently differs from either revenue or allocative efficiency. Around the world, local governments of large cities impose rent control policies and run affordable housing programs. Developed and developing countries alike provide food at subsidized prices to some populations (e.g., food stamps in the US, or in-kind provision of rice and wheat in India). Citizens of many countries can use publicly provided (but typically low-quality) basic health care. Access to roads is free (or relatively cheap), leading to inefficient levels of congestion in most large agglomerations. Covid-19 vaccines were distributed (effectively) free of charge, despite their initial scarcity. Following the spike in energy prices caused by Russia's invasion of Ukraine, virtually all European countries began subsidizing or directly controlling electricity prices faced by households. National parks tend to allocate hiking permits by lottery rather than by charging a market-clearing (or revenue-maximizing) price. The list goes on.

Author's address: piotr.dworczak@northwestern.edu. The letter summarizes the research progress in my long-term collaboration with Mohammad Akbarpour and Scott Duke Kominers on Inequality-aware Market Design. I thank Andrzej Skrzypacz for helpful comments. I gratefully acknowledge support received under the ERC Starting grant IMD-101040122. Views and opinions expressed are those of the author only and do not necessarily reflect those of the European Union or the European Research Council.

In short, policymakers often care about inequalities in the markets they control—their objective function is *redistributive*, in the sense that they have preferences over splits of surplus across market participants, and are especially concerned about the welfare of the poorest (or otherwise disadvantaged) individuals.

My impression is that economic theorists have not paid enough attention to these kinds of problems because classical economic intuition—based on the celebrated welfare theorems—suggests that policies such as in-kind transfers, subsidies, or price controls are simply a mistake. According to welfare theorems, we should “redistribute endowments,” and then let markets do their job (we would only intervene if markets are not competitive or feature externalities). However, the classical intuition is incomplete, to say the least. The second welfare theorem, in particular, breaks down under asymmetric information. In all the above examples, information available to policymakers is imperfect at best. And mechanism design provides tools to address the imperfect information problem head-on.

This is the subject of the research agenda on inequality-aware market design that I have been pursuing with Mohammad Akbarpour and Scott Duke Kominers (as well as several new co-authors more recently). We employ a mechanism-design framework to understand the optimal design of markets under imperfect information when the designer has a redistributive objective function. As it turns out, in-kind provision, subsidies, and even price controls may be part of the optimal design. But—as always—understanding whether and how we want to use these policy tools is a complex question that no single paper can address.¹

My primary goal in this letter is to introduce the readers to a simple framework of inequality-aware market design, based on our two papers Dworczak [Ⓢ] Kominers [Ⓢ] Akbarpour [2021] and Akbarpour [Ⓢ] Dworczak [Ⓢ] Kominers [2024]. The interactions of economists and computer scientists, especially recently, led to a proliferation (and popularization) of new results and approaches in mechanism design. Issues related to robustness, complexity, and computational constraints are becoming part of mainstream mechanism design *across* the two fields. My hope is that some of the same ideas applied to inequality-aware market design can generate a comprehensive understanding of when and how markets should be regulated by policymakers concerned about inequality.

2. FRAMEWORK

Let us consider a simple setting of allocating a set of homogeneous indivisible goods (e.g., identical houses) to agents differing in their privately observed values for the good and social welfare weights. Welfare weights are a classical way of capturing the designer’s preferences over how surplus is split: Higher welfare weights can be placed on agents who are poor, disadvantaged, or deserve preferential treatment for other social and moral reasons. In a reduced-form way, dispersion in welfare weights can thus capture various inequalities among market participants (under the implicit assumption that these inequalities motivate the designer’s redistributive preferences).

¹This letter is not intended to serve as a guide to the literature on the topic. Nevertheless, a shoutout to a few key contributions is in place. In particular, our work was preceded by the important papers of Weitzman [1977], Spence [1977], and Condorelli [2013].

There is a unit mass of agents, with each agent characterized by a type (r, λ) , where r is the willingness to pay (WTP) for the good and λ is the social welfare weight. The designer knows the joint distribution $F(r, \lambda)$ of WTP and welfare weights but does not observe individual realizations. The unobservability of λ is a key assumption: The designer's redistributive preferences depend on characteristics (such as wealth, income, life circumstances, and health status) that may be difficult to observe or condition on.

An agent with type (r, λ) who gets the good with probability x and pays t gets utility $rx - t$ and contributes $\lambda(rx - t)$ to the social welfare function. Quasi-linearity of preferences gives λ a nice interpretation: It is the social value of giving an agent a dollar. The social welfare function is the sum of $\lambda(rx - t)$ across all agents; additionally, the designer attaches a weight $\alpha \geq 0$ to the revenue generated by the mechanism. Since α can be flexibly specified, it is without loss of generality to assume that transfers in the mechanism are non-negative: If the designer wants to give agents a lump-sum transfer (using the revenue generated by the mechanism), it suffices to set α to be equal to the average social welfare weight.

The mechanism chosen by the designer—specifying the allocation x and transfer t for each type (r, λ) —must be individually-rational and incentive-compatible. The designer can choose a total quantity Q of goods to allocate at a total cost given by $C(Q) \geq 0$. The problem of the designer is to maximize social welfare over all feasible mechanisms:

$$\max_{\substack{x: \text{supp}(F) \rightarrow [0, 1] \\ t: \text{supp}(F) \rightarrow \mathbb{R}_+}} \int (\lambda [x(r, \lambda)r - t(r, \lambda)] + \alpha t(r, \lambda)) dF(r, \lambda) - C\left(\int x(r, \lambda) dF(r, \lambda)\right)$$

subject to

$$\begin{aligned} x(r, \lambda)r - t(r, \lambda) &\geq x(r', \lambda')r - t(r', \lambda'), & \forall (r, \lambda), (r', \lambda') \in \text{supp}(F), \\ x(r, \lambda)r - t(r, \lambda) &\geq 0, & \forall (r, \lambda) \in \text{supp}(F). \end{aligned}$$

The problem can be solved by adapting a few standard steps.²

- (1) Even though the designer would like to condition the allocation on realized welfare weights λ , she cannot; this is because λ does not affect agents' individual payoffs, and hence agents would not report λ truthfully if it led to better outcomes conditional on some reports. As long as r is continuously distributed,³ it is without loss of optimality for the designer to condition the allocation and transfers on the reported WTP only. I will abuse notation slightly and denote by $F(r)$ the marginal distribution of r . I also assume F has a continuous density f with support $[0, \bar{r}]$.⁴

²See Dworczak [Ⓔ] Kominers [Ⓔ] Akbarpour [2021] and Akbarpour [Ⓔ] Dworczak [Ⓔ] Kominers [2024] for details and missing steps.

³With a discrete type space, the designer could elicit some information about λ relying on agents' indifference; see Ostrizek and Sartori [2023].

⁴That the lower bound of WTP is 0 is not an innocuous assumption; see Akbarpour [Ⓔ] Dworczak [Ⓔ] Kominers [2024] for the analysis of the general case.

- (2) As a result, the objective function can be written as

$$\int (\lambda(r) [x(r)r - t(r)]) dF(r) + \alpha \int t(r) dF(r) - C \left(\int x(r) dF(r) \right),$$

where $\lambda(r) = \mathbb{E}[\lambda | r]$. That is, even though the designer cannot elicit λ truthfully, she can still use its *statistical correlation* with willingness to pay to inform the choice of the optimal mechanism.

- (3) Once we have a standard one-dimensional screening problem, we can use the ideas of Myerson [1981]: Incentive-compatibility requires $x(r)$ to be weakly increasing in r and pins down transfers uniquely, up to the utility $\underline{U}$ of the agent with the lowest WTP. Because I assumed that $\min(\text{supp}(F)) = 0$ and transfers are non-negative, we must in fact have $\underline{U} = 0$. Then, we know that

$$t(r) = x(r)r - \int_0^r x(\tau) d\tau, \forall r.$$

- (4) A simple sequence of standard transformations applied to the social welfare function yields that the objective function (gross of costs) is given by

$$\int (\Lambda(r)h(r) + \alpha J(r)) x(r) dF(r),$$

where $\Lambda(r) = \mathbb{E}[\lambda(\tilde{r}) | \tilde{r} \geq r]$ is the average welfare weight on agents with WTP above r , $h(r) = (1 - F(r))/f(r)$ is the inverse hazard rate, and $J(r) = r - h(r)$ is the virtual surplus function. Let $V(r) \equiv \Lambda(r)h(r) + \alpha J(r)$ and note that this value function collapses to well-known objects in special cases: willingness to pay itself if the designer has no redistributive preferences ($\lambda \equiv 1 = \alpha$); virtual surplus if the designer does not care about agent welfare ($\lambda \equiv 0$); and the inverse hazard rate in the case when the designer does not have redistributive preferences and screens with a costly-screening instrument rather than with money ($\lambda \equiv 1$ and $\alpha = 0$).

- (5) Whenever the objective function $V(r)$ fails to be non-decreasing, the constraint requiring the allocation rule $x(r)$ to be non-decreasing might bind at the optimal solution. I will therefore apply a variant of Myerson's ironing procedure. Let us first rewrite the problem in the quantile space:

$$\int_0^{\bar{r}} V(r) x(r) dF(r) = \int_0^1 V(F^{-1}(q)) x(F^{-1}(q)) dq.$$

We can optimize over non-decreasing (quantile) allocation rules $\tilde{x} : [0, 1] \rightarrow [0, 1]$ given by $\tilde{x}(q) = x(F^{-1}(q))$. Since each candidate $\tilde{x}(r)$ is a weakly-increasing function from $[0, 1]$ to $[0, 1]$, we can formally treat it as a cumulative distribution function (it is without loss of generality to assume that $x(0) = 0$ and $x(1) = 1$). Moreover, if Q is the total quantity of objects, then $1 - Q$ is the mean of the distribution $\tilde{x}$:

$$1 - Q = 1 - \int_0^{\bar{r}} x(r) dF(r) = 1 - \int_0^1 \tilde{x}(t) dt = \int_0^1 t d\tilde{x}(t).$$

An analogous transformation of the objective function (using integration by

parts) yields an equivalent representation of the problem, for a fixed Q :

$$\max_{\text{cdf } \tilde{x}} \int_0^1 \left(\int_t^1 V(F^{-1}(q)) dq \right) d\tilde{x}(t) \quad (1)$$

$$\text{subject to } 1 - Q = \int_0^1 t d\tilde{x}(t). \quad (2)$$

Thus, for a fixed quantity Q , the problem is to choose a distribution $\tilde{x}$ to maximize the objective function subject to preserving a given mean of the distribution.⁵

- (6) Let $\mathcal{V}(t) = \int_t^1 V(F^{-1}(q)) dq$, and let $\text{co}\mathcal{V}$ denote the concave closure of $\mathcal{V}$ (the smallest concave function that lies above $\mathcal{V}$). The value of problem (1)-(2) is $\text{co}\mathcal{V}(1 - Q)$. Moreover, there always exists an optimal distribution $\tilde{x}$ that has at most two points in its support, which implies that the optimal allocation rule $x(r)$ takes the form

$$x(r) = \begin{cases} 0 & r < r_0, \\ x_0 & r \in [r_0, r_1], \\ 1 & r \geq r_1, \end{cases}$$

for some $0 \leq r_0 \leq r_1 \leq 1$, and $x_0 \in (0, 1)$. In particular, a posted-price mechanism is optimal if and only if the optimal $\tilde{x}$ is a degenerate distribution, which is the case if and only if $\text{co}\mathcal{V}(1 - Q) = \mathcal{V}(1 - Q)$. Rationing appears in the optimal mechanism if and only if $\text{co}\mathcal{V}(1 - Q) > \mathcal{V}(1 - Q)$, that is, whenever ironing of the objective function is required given the chosen quantity Q .

- (7) To solve the final problem, all we have to do is maximize the total objective $\text{co}\mathcal{V}(1 - Q) - C(Q)$ over possible quantity levels $Q \in [0, 1]$.

3. ECONOMIC IMPLICATIONS

For the remainder of the note, I will explore the economic consequences of the derivation. To focus on the role of welfare weights, I assume that $F(r)$ is the uniform distribution on $[0, 1]$. Let us also assume that $\alpha = \mathbb{E}[\lambda] = 1$, meaning that all revenue is redistributed back to agents as a lump-sum transfer. Under these assumptions, we obtain

$$V(r) = \int_r^1 \lambda(\tau) d\tau + 2r - 1. \quad (3)$$

To gain some intuition for the shape of the function $\lambda(r)$, recall that $\lambda(r)$ is the conditional expectation of the unobserved welfare weight λ conditional on the observed WTP r . WTP is shaped by both the taste for the good and the “opportunity cost of money”—it is natural to expect that wealthier agents will be able to

⁵This problem is mathematically equivalent to a Bayesian persuasion problem with a binary state: Choose a distribution of posterior beliefs $\tilde{x}$ to maximize an objective function (1) that depends on the posterior belief, subject to a Bayes-plausibility constraint (2). By Aumann and Maschler [1995] and Kamenica and Gentzkow [2011], the value of such a problem is equal to the concave closure of the objective function at the prior (here, $1 - Q$), and the optimal distribution has at most two points in its support.

pay more for things they need. Thus, as long as the designer has a preference for redistribution towards agents with a higher opportunity cost of money, we might expect a negative correlation of λ and r , resulting in a decreasing function $\lambda(r)$. The strength of the correlation depends on whether heterogeneity in WTP is driven more by heterogeneity in tastes or by heterogeneity in ability to pay (the negative correlation is perhaps strongest for expensive goods and services satisfying basic needs, such as health care or housing).

When is it optimal to use a price mechanism?

When inequality—as measured by the dispersion in welfare weights—is not too high, we should use a price mechanism:

PROPOSITION 3.1. *If $\lambda(r) \leq 2$ for all r , then it is optimal to use a price mechanism.*

Proposition 3.1 follows from noticing that when the expected welfare weights never exceed the average welfare weight (normalized to 1) by more than a factor of 2,⁶ then the objective function $V(r)$ is non-decreasing—ironing is not required, so a posted-price mechanism is optimal. Nevertheless, the price in the mechanism will typically be “distorted” relative to the efficient mechanism. To see that, suppose that $C(Q) = cQ$, so that the good can be produced at constant marginal cost $c > 0$. If all $\lambda \equiv 1$, it is optimal to set the price equal to marginal cost. This is not the case when the designer has redistributive preferences:

PROPOSITION 3.2. *Under the assumptions of Proposition 3.1, if $C(Q) = cQ$ for $c > 0$, the optimal price p^* satisfies*

$$p^* = c + (1 - \Lambda(p^*))(1 - p^*).$$

It follows that the designer optimally imposes a tax (the price p^* is higher than the marginal cost c) if and only if the average welfare weight $\Lambda(p^*)$ on agents who buy the good is lower than the average welfare weight (normalized to 1). This allows the designer to increase the revenue and hence the lump-sum transfer. This case arises when $\lambda(r)$ is decreasing in r . Otherwise, when $\Lambda(p^*) > 1$, the designer uses a subsidy ($p^* < c$): This allows the designer to transfer utility from the average agent to an agent who buys the good. Only in the knife-edge case $\Lambda(p^*) = 1$ is the efficient mechanism (pricing at marginal cost) optimal.

When is it optimal to ration?

If inequality in the market is large, it may be optimal to complement a posted price with a lower price but subject to rationing ($x_0 < 1$):

PROPOSITION 3.3. *Suppose that $\lambda(r) > 2$ for some r . Then, there exists a cost function $C(Q)$ such that the optimal mechanism features rationing.*

The assumption that $\lambda(r) > 2$ for some r implies that the objective function $V(r)$ is decreasing in some region, and hence ironing may be needed (ironing is used if the

⁶The reader may reasonably ask: “Why 2?” There is a good reason which I explain in Dworczak [2023] (see also the Online Appendix of Dworczak [Ⓢ] Kominers [Ⓢ] Akbarpour [2021]).

cost function makes the optimal quantity Q fall within the ironing region). In our context, ironing means that the optimal mechanism offers a rationed option with a lower per-unit price. The role of rationing is to shift more gains towards agents with intermediate values of r : Agents with the highest r self-select into the non-rationed high-price option, which allows the designer to subsidize agents choosing the rationed option by reducing its price. However, because rationing is inefficient, the welfare weights on agents who benefit from it must be sufficiently large.

At least among economists, Myerson’s ironing is mostly perceived as an “annoying detail” that can be formally ignored if we make appropriate regularity conditions. With redistributive preferences, however, ironing is no longer a technicality. It underlies the key economic result: Inefficient rationing may be part of the optimal mechanism.

The assumption that the expected welfare weight $\lambda(r)$ exceeds 2 for some willingness to pay r has an important economic interpretation. It states that the designer has strong redistributive preferences to begin with (a necessary condition, obviously, is that there are at least some agents with $\lambda > 2$) *and* that market behavior reveals information about that underlying inequality (i.e., r is sufficiently informative about λ). Intuitively, even if the underlying inequality is substantial, rationing is suboptimal if willingness to pay is driven more by tastes than ability to pay.

4. DIRECTIONS FOR FUTURE RESEARCH

Relative to the simple framework presented here, recent work incorporated several realistic features such as (i) two-sided markets with both buyers and sellers (Dworczak [Ⓔ] Kominers [Ⓔ] Akbarpour [2021]), (ii) heterogeneous object qualities and partially informative signals available to the designer (Akbarpour [Ⓔ] Dworczak [Ⓔ] Kominers [2024]), (iii) allocative externalities (Kang [2024], Pai and Strack [2022], Akbarpour [Ⓔ] Budish [Ⓔ] Dworczak [Ⓔ] Kominers [2024]), (iv) co-existence of regulated and private markets (Kang [2023]), (v) endogeneity of buyer and seller roles (Kang and Zheng [2023]), (vi) models with finitely many agents (Reuter and Groh [2020]), (vii) the role of costly screening devices (Dworczak [2023], Yang et al. [2024]), and (viii) the use of information design and market segmentation (Barreto et al. [2022], Arya and Malhotra [2022]). However, all of these papers adopt a standard Bayesian optimization framework. I conclude this letter by identifying a few research directions along which the expertise of the Econ CS community would be particularly valuable.

- (1) **Price of anarchy.** A classical question asked in the computer science literature concerns the performance of a system relative to its optimal design, evaluated against the worst-case scenario (see, e.g., Koutsoupias and Papadimitriou [2009]; Roughgarden and Tardos [2000]; Andelman et al. [2009]). One can ask a version of that question when performance (e.g., of a market) is measured in an inequality-sensitive way, for example, using a redistributive objective function. Just how bad can an unregulated market be in achieving an equitable outcome? Are there circumstances under which we should be particularly concerned about inequality in markets?

- (2) **Performance of simple mechanisms.** A closely related question is that of the performance of simple (non-optimal) mechanisms (in the context of allocating goods, see, e.g., Hartline and Roughgarden [2009]; Chawla et al. [2010]; Feldman et al. [2015]; Babaioff et al. [2020]; Feldman et al. [2020]). In the above framework, the optimal mechanism is relatively simple, in that it involves at most two prices (one at which agents can buy for sure, and one with rationing). Even there, however, one can ask: what fraction of the optimal welfare can the designer achieve by relying only on posted-price mechanisms? In other words, is the use of rationing essential to achieving a good welfare guarantee? In more complex settings, the optimal mechanism can become very complicated, and finding simpler mechanisms that achieve good welfare guarantees (in terms of the redistributive objective function) is of first-order importance for practical purposes.
- (3) **Robustness.** The computer scientists' insistence on robustness has undoubtedly (at least partially) inspired the distributionally-robust mechanism design literature that is flourishing in economics (see, e.g., Frankel [2014], Carroll [2017], Brooks and Du [2021], Suzdaltsev [2022], He and Li [2022], among many others). Whether rationing is used in the above framework depends on the joint distribution of willingness to pay and welfare weights: r and λ must be sufficiently strongly correlated. But what if the designer does not know their exact joint distribution? What if she considers a robust approach of maximizing her (redistributive) objective against the worst-case joint distribution consistent with some known moments?
- (4) **Multidimensional screening.** One of the most spectacular advancements brought together by the interaction of economics and computer science is in the area of multi-dimensional screening (see, e.g., Daskalakis et al. [2013, 2017]; Haghpasand and Hartline [2020]). The framework I considered here features two-dimensional types but the underlying screening problem ends up being one-dimensional. Many natural extensions of the framework, however, feature multi-dimensional information that cannot be reduced to a one-dimensional statistic. For example, the designer could be controlling multiple markets, or have access to multiple redistributive instruments. The question of optimal design with an inequality-sensitive objective in such environments is wide open.

Most likely, the best future research directions are not on the above list. The beauty of good ideas is that they cannot be anticipated. And we need more good ideas to provide smart policy guidance for dealing with inequality (policymakers will try to deal with it, with or without our input!).

REFERENCES

- AKBARPOUR, M. [Ⓢ] E. BUDISH [Ⓢ] P. DWORCZAK [Ⓢ] S. D. KOMINERS. 2024. An economic framework for vaccine prioritization. *Quarterly Journal of Economics* 139, 1, 359–417.
- AKBARPOUR, M. [Ⓢ] P. DWORCZAK [Ⓢ] S. D. KOMINERS. 2024. Redistributive allocation mechanisms. *Journal of Political Economy* (forthcoming).
- ANDELMAN, N., FELDMAN, M., AND MANSOUR, Y. 2009. Strong price of anarchy. *Games and Economic Behavior* 65, 2, 289–317.
- ARYA, Y. AND MALHOTRA, R. 2022. Rethinking distribution: Introducing market segmentation as a policy instrument. Working Paper.

- AUMANN, R. J. AND MASCHLER, M. 1995. *Repeated Games with Incomplete Information*. MIT press.
- BABAIOFF, M., IMMORLICA, N., LUCIER, B., AND WEINBERG, S. M. 2020. A simple and approximately optimal mechanism for an additive buyer. *J. ACM* 67, 4 (jun).
- BARRETO, D. M. A., GHERSENGORIN, A., AND AUGIAS, V. 2022. Price discrimination with redistributive concerns. Working Paper.
- BROOKS, B. AND DU, S. 2021. Optimal auction design with common values: An informationally robust approach. *Econometrica* 89, 3, 1313–1360.
- CARROLL, G. 2017. Robustness and separation in multidimensional screening. *Econometrica* 85, 2, 453–488.
- CHAWLA, S., HARTLINE, J. D., MALEC, D. L., AND SIVAN, B. 2010. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the Forty-Second ACM Symposium on Theory of Computing*. STOC '10. Association for Computing Machinery, New York, NY, USA, 311–320.
- CLARKE, E. H. 1971. Multipart pricing of public goods. *Public Choice* 11, 1, 17–33.
- CONDORELLI, D. 2013. Market and non-market mechanisms for the optimal allocation of scarce resources. *Games and Economic Behavior* 82, 582–591.
- DASKALAKIS, C., DECKELBAUM, A., AND TZAMOS, C. 2013. Mechanism design via optimal transport. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*. EC '13. Association for Computing Machinery, New York, NY, USA, 269–286.
- DASKALAKIS, C., DECKELBAUM, A., AND TZAMOS, C. 2017. Strong duality for a multiple-good monopolist. *Econometrica* 85, 3, 735–767.
- DWORCZAK, P. 2023. Equity-efficiency trade-off in quasi-linear environments. Working Paper.
- DWORCZAK, P. [Ⓢ] S. D. KOMINERS [Ⓢ] M. AKBARPOUR. 2021. Redistribution through markets. *Econometrica* 89, 4, 1665–1698.
- FELDMAN, M., FU, H., GRAVIN, N., AND LUCIER, B. 2020. Simultaneous auctions without complements are (almost) efficient. *Games and Economic Behavior* 123, 327–341.
- FELDMAN, M., GRAVIN, N., AND LUCIER, B. 2015. Combinatorial auctions via posted prices. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA '15. Society for Industrial and Applied Mathematics, USA, 123–135.
- FRANKEL, A. 2014. Aligned delegation. *American Economic Review* 104, 1 (January), 66–83.
- GROVES, T. 1973. Incentives in teams. *Econometrica* 41, 4, 617–631.
- HAGHPANAH, N. AND HARTLINE, J. 2020. When is pure bundling optimal? *The Review of Economic Studies* 88, 3 (08), 1127–1156.
- HARTLINE, J. D. AND ROUGHGARDEN, T. 2009. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM Conference on Electronic Commerce*. EC '09. Association for Computing Machinery, New York, NY, USA, 225–234.
- HE, W. AND LI, J. 2022. Correlation-robust auction design. *Journal of Economic Theory* 200, 105403.
- KAMENICA, E. AND GENTZKOW, M. 2011. Bayesian persuasion. *American Economic Review* 101, 6, 2590–2615.
- KANG, M. AND ZHENG, C. Z. 2023. Optimal design for redistributions among endogenous buyers and sellers. *Economic Theory*, 1141–1180.
- KANG, Z. Y. 2023. The public option and optimal redistribution. Working Paper.
- KANG, Z. Y. 2024. Optimal indirect regulation of externalities. Working Paper.
- KOUTSOUPIAS, E. AND PAPADIMITRIOU, C. 2009. Worst-case equilibria. *Computer Science Review* 3, 2, 65–69.
- MYERSON, R. 1981. Optimal auction design. *Mathematics of Operations Research* 6, 1, 58–78.
- OSTRIZEK, F. AND SARTORI, E. 2023. Screening while controlling an externality. *Games and Economic Behavior* 139, 26–55.
- PAI, M. AND STRACK, P. 2022. Taxing externalities without hurting the poor. Working Paper.

- REUTER, M. AND GROH, C.-C. 2020. Mechanism design for unequal societies. *Working Paper*.
- ROUGHGARDEN, T. AND TARDOS, É. 2000. How bad is selfish routing? *j acm. Foundations of Computer Science, 1975., 16th Annual Symposium on* 49, 93–102.
- SPENCE, M. 1977. Nonlinear prices and welfare. *Journal of Public Economics* 8, 1, 1–18.
- SUZDAL'TSEV, A. 2022. Distributionally robust pricing in independent private value auctions. *Journal of Economic Theory* 206, 105555.
- VICKREY, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* 16, 1, 8–37.
- WEITZMAN, M. L. 1977. Is the price system or rationing more effective in getting a commodity to those who need it most? *Bell Journal of Economics* 8, 2, 517–524.
- YANG, F., DWORCZAK, P., AND AKBARPOUR, M. 2024. Comparison of screening devices. *Working Paper*.

Generative AI as Economic Agents

NICOLE IMMORLICA and BRENDAN LUCIER

Microsoft Research New England

and

ALEKSANDRS SLIVKINS

Microsoft Research New York City

Traditionally, AI has been modeled within economics as a technology that impacts payoffs by reducing costs or refining information for human agents. Our position is that, in light of recent advances in generative AI, it is increasingly useful to model AI itself as an economic agent. In our framework, each user is augmented with an AI agent and can consult the AI prior to taking actions in a game. The AI agent and the user have potentially different information and preferences over the communication, which can result in equilibria that are qualitatively different than in settings without AI.

Categories and Subject Descriptors: J.4 [Social and Behavioral Science]: Economics

General Terms: Economics, Theory

Additional Key Words and Phrases: Generative AI, Economic Agents

In economic theory models, agents take potentially costly actions to maximize their expected utility given their information about the world and others' strategies. In some models, agents have access to technologies that help by reducing action costs or providing informative signals. Researchers have tended to model AI as such a technology. However, given the recent advancements in generative AI, and particularly Large Language Models (LLMs), we claim this new technology is itself sometimes best modeled as an economic agent.

Why do we suggest modeling generative AI as an agent in a game? After all, it is just a type of machine-learning model that predicts the next token in a sequence or samples content from a learned distribution. A notable aspect of generative AI, however, is its ability to generate novel content based on an implicit but vast "common-sense" understanding of the world.¹ As such, generative AI models (and systems built upon them) can be leveraged as virtual consultants that assist, analyze, or even strategize on behalf of their users. We argue below that, much like a human consultant, an AI-powered virtual consultant exhibits features typical of an economic agent.

Indeed, like a human consultant, an AI consultant has an *information set*, or a collection of knowledge and beliefs about the world and other actors therein.

¹While such anthropomorphic metaphors can be misleading, our argument requires only a *semblance* of common-sense reasoning.

Authors' addresses: {nicimm, brlucier, slivkins}@microsoft.com

We would like to thank Nageeb Ali, Drew Fudenberg, John Horton, Adam Kalai, David Kempe, Akshay Krishnamurthy, Kevin Leyton-Brown, Yishay Mansour, Sam Taggart, and Clayton Thomas for their many helpful comments.

This information can come from the training data, prompts, or external sources. Trained on massive datasets, AI consultants may have different, possibly much richer, signals about the world than any human. Moreover, the potential to learn from interactions with their entire user base opens up opportunities to adapt and integrate new knowledge more quickly than a human consultant.

An AI consultant also has an *action space* consisting of outputs to queries. Similar to a human consultant, this action space can be very general, including, e.g., natural language or source code. However, current social norms impose constraints on the environment in which the AI can take these actions. Namely, the AI is constrained to a virtual environment: its communication with a user. Thus, AI has *limited agency*; while it has agency over its responses, the user retains the power to make (or, equivalently, to veto) any payoff-relevant actions, possibly based on its interactions with the AI.²

Finally, an AI consultant has objectives and constraints ingrained during training, fine-tuning, and orchestration. This induces the AI to act as though it is maximizing some implicit *preferences*.³ Importantly, similar to human consultants, these preferences may not be perfectly aligned with the user’s preferences and are not directly under the user’s control. However, unlike human consultants, the AI, once deployed, has a limited view of the world that typically consists only of communications with its user. Thus its preferences can only be a function of this communication transcript rather than real-world outcomes.⁴

Thus, our view is that AI-based virtual consultants can be reasonably modeled as economic agents. As discussed above, they have limited agency relative to human consultants, as well as limited worldview and preferences, but can potentially be more knowledgeable and adaptive. This view of AI as an economic agent includes, as a special case, classic models that treat AI as a technology that directly reduces action costs or provides informative signals, as these can be viewed as agents with trivial preferences and particularly structured information and action sets.

AI consultants can also be drastically cheaper than human consultants, and hence applicable in a wider array of contexts. Across these contexts, they can play three (overlapping) roles: *assistant*, completing specific tasks; *analyst*, revealing and conveying information from built-in knowledge or external sources; and *strategist*, proposing actions for a user and possibly accounting for other parties’ reactions. The three roles can lead to different modeling choices from an economic perspective.

Synthesizing these considerations, we propose augmenting economic theory models so that traditional agents (henceforth *users*) are endowed with AI-based con-

²We model AI in this way to reflect our focus on its role as a *consultant*. Our model also applies when AIs can directly take actions in the real world, as long as the user retains veto power over such actions.

³We distinguish between “inherent” preferences of the AI and those induced by user prompts (e.g., “you are a candy-loving baby”). Note that the former can shape the latter by, e.g., influencing the manner and extent to which prompts are followed.

⁴For example, a human consultant paid to work on one part of a larger project might prefer that the project is ultimately successful because this would look good on a resume. In contrast, an AI consultant in our model has no awareness of (and thus no incentives regarding) a project’s real-world outcome beyond the user’s instructions.

sultants (henceforth *AI agents* or simply *AIs*) and can interact with the latter before taking actions. In what follows, we put forward a formal framework to study the impact of AI agents as a step toward informing their design and deployment. We instantiate this framework with several examples (including examples where preferences are misaligned), outline potential research questions, and relate our framework to existing literatures in economics and AI. Our view is that tools and techniques from the theory of economics and computation are perfectly suited to the challenges that arise in a world of AI agents, where the inherent features of AI systems meet the game-theoretic implications of individual preferences and multi-agent interactions.

1. A GENERAL MODEL

Baseline Model: Single-Shot Game without AI Agents. The baseline for our model is the following general game-theoretic setup. There are n (human) players who interact in a game. Each player i has a (finite or infinite) action space A_i . An action profile $\vec{a} \in A := \prod_i A_i$ denotes a choice of action by each player. There is also a special agent called Nature who selects a state of nature $\omega \in \Omega$ at random from a commonly-known distribution. Each player i is endowed with an information structure: a partition $\mathcal{I}_i$ of Ω .⁵

The game proceeds as follows. Nature first selects a state of nature ω . Each player i is then endowed with an information set $I_i \subseteq \Omega$ such that $\omega \in I_i \in \mathcal{I}_i$, which describes her (partial) information about the chosen state of nature. Players then simultaneously select actions. Finally, each player receives a payoff $u_i(\vec{a}, \omega)$ that depends on the actions of all players and the state of nature.

This model can capture multiple game-theoretic scenarios and could be paired with various solution concepts. If there is only a single state of nature, this is a standard simultaneous-move game, and one might consider pure or mixed Nash equilibria. To give another example, the state of nature might consist of a vector of types, one per human player, and Nature might draw this state from a known product distribution. Each player's information set might then reveal their own type, but not that of the other players. Then we are in a Bayesian model of incomplete information, and we may be interested in its Bayes-Nash equilibrium. Our framework can also be applied to other game formats such as extensive form games in which the information sets of the players contain (possibly incomplete) information about the history of play.

AI Agents. We augment our baseline model by associating an AI agent with each human player. The human and the AI agent interact through a communication protocol, in which they alternate sending messages from a space M of potential messages beginning with the human player. Communication can proceed in multiple rounds, resulting in a *transcript* τ_i (i.e., an ordered sequence of messages) between human player i and their AI agent. Our only assumption about the messaging protocol is that the human player can always choose not to send a message, effectively terminating the communication procedure. Furthermore, and central to our work, the AI agent's payoff might be misaligned with that of its human player.

⁵That is, $\mathcal{I}_i$ is a collection of disjoint subsets whose union is Ω .

Like the human players, each AI agent i has an information structure $\mathcal{J}_i$ and an information set $J_i \subset \Omega$ such that $\omega \in J_i \in \mathcal{J}_i$ that specifies what they know about the state of nature. Each AI agent i also has a payoff $v_i(\tau_i, \omega)$ that can depend on the state of nature and their transcript, but is otherwise independent of the actions chosen by the human players. The information set held by an AI agent is not necessarily the same as the information set of their human user, and indeed a key aspect of AI agents is that they may have access to information that their associated human player does not.

We incorporate AI agents into the timing of our general game as follows. Nature first selects a state $\omega \in \Omega$, as before. Each human player i then engages in communication with their associated AI agent, generating a (possibly empty) message transcript τ_i . There is a cost function c_i that associates with each message transcript τ_i a cost $c_i(\tau_i) \geq 0$, normalized so that an empty transcript costs 0. Once communication is complete, each human player i then selects an action $a_i \in A_i$ as before. The payoff to each human player is $u_i(\vec{a}, \omega) - c_i(\tau_i)$, their payoff from the baseline game less any costs incurred from communication with the AI agent. The AI agent receives payoff $v_i(\tau_i, \omega)$.

We emphasize that the payoff to an AI agent can depend on τ_i and ω , but does not directly depend on realized outcomes or payoffs from the baseline game. Consequently, the AI agent's payoff is determined once its interaction with the user is complete. In essence, a virtual consultant is evaluated on how well it provides advice to its user, rather than how the user chooses to act on that advice. For example, the AI agent's payoff might depend on the accuracy of information it provides, or even on its user's *anticipated* utility from the baseline game under some behavioral model for the players. But we do not allow a direct dependence on the user's *realized* utility, as this falls outside the scope of the AI agent's interaction with the user.

2. EXAMPLE SCENARIOS

Before we explore ways that an AI agent's misaligned preferences can impact outcomes, recall that our framework also encompasses settings where an AI agent is simply a technology that reduces costs or increases information by providing signals. Even such simple settings lead to counter-intuitive outcomes: for example, as with other cost-reducing or signal-inducing technologies, the presence of an AI technology can modify a user's choice set in ways that ultimately reduce the payoffs of the human agent. See Appendices A.1-A.2 for concrete examples of this dilemma in the context of an email game.

That said, our primary focus is on scenarios where the AI is an agent with preferences that may not be not fully aligned with the user. We note that while many are concerned with AIs having malicious intent, misalignment can cause harm even when the AI is essentially benign.⁶ The following examples illustrate some of the ways in which our model captures the impact of AI preferences on outcomes.

⁶More generally, misalignment can have non-trivial game-theoretic implications that can be either positive or negative. We focus on harms only for ease of presentation.

2.1 Example 1: Advice Evaluated on Perceived Helpfulness

A single user is deciding between a set A of possible actions and can consult an AI agent with access to payoff-relevant information $\omega \in \Omega$. For instance, the user might be deciding between two similar products and the AI agent knows their features. The user’s preferences are described by a utility function $u: \Omega \times A \rightarrow \mathbb{R}$. A typical interaction may involve the user describing their preferences and the AI agent responding with relevant information or recommendations. In such a “chatbot” scenario, an AI is typically evaluated on (and hence optimized for) metrics like correctness, perceived helpfulness, and relevance. These goals, which we can think of as an AI’s preferences, seem aligned with the user at first glance: after all, helpful, relevant, and personalized advice should lead to better decisions. But we emphasize that these metrics are functions only of the state of the world and the communication transcript, not the user’s realized utility. Essentially, the AI agent maximizes *perceived helpfulness* at the moment of communication, rather than realized helpfulness in retrospect.

While there may be different ways to operationalize “perceived helpfulness,” one can imagine that many of these lead to an AI agent being over-rewarded for being convincing and/or suggesting a new course of action, rather than providing a more balanced view. In particular, the AI may over-emphasize unexpectedly positive or negative features of certain choices, even if they are not fully representative. Such misalignment may be problematic even if the AI agent is constrained to only return factual information.⁷

We provide a formal numerical example in Appendix A.3, demonstrating that this effect can cause a rational user to derive no benefit from the AI agent even though it only ever provides accurate information. Of course, there may be other ways to set the AI agent’s preferences that would avoid this particular issue, such as an optimization loop in which user satisfaction is judged only in hindsight after actions are taken. Our point is that the way the AI agent’s payoff function is defined can generate unintended distortions by misaligning incentives, and optimizing for a user’s declared satisfaction at the moment of interaction with the AI agent could be one such source of distortion.

2.2 Example 2: Delegated Search

A human user is attempting to optimize over a space Γ of options (think documents or images), and can delegate this task to an AI agent. The user has a utility function $u: \Gamma \rightarrow \mathbb{R}$ over the space Γ of all *conceivable* options, from some class U of possible utility functions. However, only a subset $X \subseteq \Gamma$ of the options are actually implementable. For example, Γ might consist of all possible sequences of English words and punctuation, whereas X is a set of grammatically correct and meaningful documents. The utility function u and feasible set X constitute the state of nature, drawn from a known prior. We posit information asymmetry: the user knows u but not X , and the AI agent knows X but not u .⁸ The user can communicate a

⁷It could also incentivize the AI to hallucinate, but we do not model that in this example.

⁸When we say that the user does not know X , what we require is only that they cannot perform optimization over the (possibly very complex) manifold of options. We allow that the user might recognize an element of X “when they see it.”

utility function $u' \in U$ to the AI agent, and the latter returns an option $x \in X$. The user's payoff is $u(x)$. The AI agent's payoff is $u'(x) - \gamma(x)$, where γ is a fixed penalty function (which can be infinite) chosen by the agent's designer to steer the agent away from some problematic outcomes.

When X is known to the user (i.e., it is drawn from a point mass distribution), this setup defines a Stackelberg game with partially aligned preferences. Each choice of $u \in U$ has a corresponding maximizer $x(u) \in X$ for the AI agent. We should expect the human user to report $u' \in U$ that maximizes $u(x(u'))$, yielding the Stackelberg payoff of the game. That is, the user strategically adapts her stated preferences, anticipating the influence of the penalty function.⁹

More generally, the user only knows a prior distribution over X . Then the user may still be incentivized to misreport u , just like the full-information case, but this may now lead to low-utility outcomes on some realizations of X . We provide a numerical example in Appendix A.4, illustrating that attempts to guide an AI agent toward favorable outcomes can ultimately lead to a loss of value when the user is uncertain about the space of available options.

3. RESEARCH QUESTIONS

We have described how to view AI-based virtual consultants as economic agents and integrate them into some examples of decision- and game-theoretic scenarios. For any given scenario of interest, the resulting economic models can be studied to shed light on key questions about the impact and design of AI agents. These research questions span both economics and computer science, as outlined below.

Equilibrium Analysis. Suppose first that we treat the design of an AI agent as exogenous and fixed. How does the introduction of the resulting virtual consultants influence strategic behavior and outcomes at equilibrium, both immediately and in the long term? Do these AI agents increase aggregate welfare for users in the system, and what are the distributive effects? Which users stand the benefit the most, and which are worse off? What are the implications for fairness, and are existing biases amplified or suppressed?

While the addition of AI agents can substantially shift the set of equilibria of behavior, there is also the potential to influence the choice of equilibria in games where the equilibrium is not unique. For example, reliance on AI agents might impact the coordination or correlation of actions of users or shift norms in large populations. Do cost reductions make it more difficult to build and maintain reputation? Does the ability to delegate effort to an AI agent make it more or less credible to commit to a given course of action?

Market Design. One can also take a more prescriptive approach, and consider the design of virtual consultants and/or platforms to shape incentives and influence outcomes. The design space for virtual consultants includes both the communication protocol for interfacing with the user, as well as the preference function that guides

⁹If γ and X are perfectly known and the set U includes all possible functions mapping Γ to $\mathbb{R}$, then choosing $u' = u + \gamma$ would always maximize $u(x(u'))$ and, effectively, align the AI with user's incentives. But this generic approach is not necessarily feasible if one or more of these assumptions are relaxed; we describe one such example in Appendix A.4.

the AI agent’s choice of responses. Given a game-theoretic scenario of interest, how should these be designed to improve the quality of outcomes when employed by (strategic) human users? When multiple AI agents are involved, our goal might be to design one of the AI agents fixing those available to other users, or we may be jointly designing all (or some) of the AI agents to achieve good aggregate outcomes.

One can also explore platform and market design under the assumption that users have access to AI agents. In the context of our general model, this corresponds to designing (some aspects of) the game played between the human users. As the adoption of virtual consultants becomes more widespread, how does this influence user behavior, and how should platforms and markets adapt? Should platforms provide their own AI agents, and is there a benefit to designing market rules and AI agents jointly? What sorts of emergent behaviors might arise from the resulting interactions, and to what extent might optimal market designs make AI agent usage mandatory (or disadvantage those without access)?

Algorithm design. Interactions with virtual consultants involve *algorithms*, be they in the implementation of the AI agents themselves or in the way they are used. Implicit in all of our models is the question of optimizing one’s communication with an AI agent, whether directly via algorithmic layer or indirectly through human user behavior. Likewise, an AI agent faces an optimization problem when generating responses, which can be guided by explicit algorithmic architectures built on top of LLMs. Which algorithm designs, on either side, lead to more desirable outcomes? How should performance be evaluated, and how should we value robustness and stability versus expected performance? How should the cost of using an AI agent be modeled, and how can we optimize performance-cost tradeoffs?

4. EXTENDING THE GENERAL MODEL

We presented a model that aims to capture the introduction of AI-powered virtual consultants as economic agents with limited agency. There are natural ways to extend this perspective to encompass other economic questions related to AI agents. Here we briefly describe a few high-level questions that one might extend our general model to include.

Increased AI Agency. We intentionally modeled AI agents as consultants with a limited degree of agency. In our model, AI agents can send messages to their human users but cannot (a) interact directly with other virtual consultants or other human users, or (b) directly take actions that are payoff-relevant for human users. As AI technology evolves, we may see scenarios in which AI agents become empowered to do one or both of these. For example, one might imagine an AI agent for resume screening that is empowered to reject candidates without human input in certain cases. From a modeling perspective, one might still attempt to model such a scenario as a human-taken action, implemented through a sign-off policy that the human user ultimately has control over. However, if delegating actions to AI agents becomes increasingly more commonplace, we may reach a point where it is more natural to endow agents with increased (but still limited) agency within our models.

Platforms and Interfaces for AI Agents. Our model treats an AI agent’s information set and internal optimization as intrinsic to the agent. However, one might explicitly model the manner in which an agent acquires external data or otherwise interacts with the world, especially if one also increases the agency of an AI agent. For example, an AI agent that provides recommendations for jobs to apply to (or one that helps screen job candidates) might benefit from interacting with a job-search platform. But if such agents become commonplace, it is natural to imagine that a platform could be incentivized to offer data interfaces that are tailored to virtual consultants. More broadly, one can imagine ecosystems for services and/or data to be used by virtual consultants, enabling general-purpose personal AI assistants that can serve as interfaces with the broader world through one’s smartphone. And indeed some nascent markets for virtual consultant services are already emerging.¹⁰ Modeling the ways in which virtual agents can interact with platforms, and/or how any platform usage costs would be internalized by agents (or otherwise how platforms might monetize such interactions), is likely to be of growing relevance as these ecosystems mature.

Economics of Training AI Agents. Our modeling to this point treats AI agents as entities in a scenario, but one might additionally wish to model the manner in which such agents are created and/or maintained. This might involve modeling the data used to train ML models that drive the AI agents, the way that an AI agent changes in response to its training data, and the firm(s) that provide these AI agents to users. By modeling the source of the data used to train AI agents (which may come directly from the real world, or may be generated by another AI), one can explore the economics of data generation and monetization. How are content creators incentivized to contribute to the training data, and how does this influence the content they create? How are the distribution platforms or other repositories on which that data resides incentivized to monetize and/or protect that data, and how does this impact the data on which AI agents are trained? What influence might this have on the behavior, information sets, or (implicit) incentives of AI agents?

5. RELATED LITERATURES

Our framework is related to several literatures in economics and computer science. We highlight these connections (and outline the differences) in what follows.¹¹

Delegation and Contracting. Our central argument is that modern AI assistance technologies are best modeled as economic agents. As such, our framework is closely related to principal-agent models in which a principal decision maker delegates to another (human) agent. This includes the theory of delegation and of contract design, both well-studied in microeconomics (see, e.g., [Mas-Colell et al. 1995, Chapter 14] and [Laffont and Martimort 2002] for an introduction to principal-agent models, [Bendor et al. 2001] for an overview of the theory of delegation, and [Holmstrom 1984; Armstrong and Vickers 2010] for other representative delegation mod-

¹⁰E.g., plugins for use with OpenAI’s ChatGPT.

¹¹Throughout, we only provide brief, exemplary citations, opting for surveys whenever possible. A more detailed survey of these literatures is beyond our scope.

els) and increasingly in economics and computation.

The theory of delegation focuses largely on consequences of misaligned incentives, where the principal and agent may prefer different actions be taken. The principal can attempt to mitigate this issue by committing to the way in which they will act on advice, or restricting the space of options available to the agent. A crucial aspect of such models is that the agent has preferences over the choice of action that is ultimately taken. In contrast, we model an AI agent as having preferences over the communication transcript itself, rather than actions or realized outcomes. This might mean, for example, that the agent is indifferent between aspects of the principal's preferences that are not revealed during communication, even if they would influence the principal's choice of action. In this sense our modeling of AI agent incentives is more closely related to contract design, in which agents commonly have preferences driven by the cost of taking actions but not over the outcome realizations. One difference between our model and contract design is that we assume AI agents cannot be incentivized with monetary transfers; rather, they are internally incentivized with preferences over the communication process, as mentioned above.

Algorithmic Interfaces and Human-AI Collaboration. It is increasingly common for online platforms and markets to support AI-powered interfaces and algorithms that offload platform-specific actions. These include autobidding services popularized in advertising platforms [Google 2023; Microsoft 2023; Meta 2023], price prediction and recommendation services for matching platforms like AirBnB and Amazon [Airbnb 2023; Amazon 2023], high-frequency trading algorithms [O'hara 2015], and more. There is a substantial and growing literature on the design and use of such systems, and how they ultimately impact the behavior of (human) users. Of particular interest is the goal of automating the behavior of the users themselves, as a way of performing those tasks more quickly, cheaply, and/or responsively than a human can. Recent work has begun to explore the use of generative AI in such scenarios, including automating the generation of content, and implications for platform design [Duetting et al. 2023; Yao et al. 2024; Fish et al. 2024]. Relative to these scenarios, the AI tools we seek to model are less tied to the structure of a specific domain or platform and, crucially, do not directly take actions on their users' behalf.

The design of AI-powered systems and interfaces is heavily informed by the theory of Human-AI collaboration and user behavior. This literature focuses on human-in-the-loop systems, analyzing the performance of AI systems in the context of (and in anticipation of) how they will be used [Grosz 1996; Amershi et al. 2019]. This includes analyzing the importance of features like trust [Hoff and Bashir 2015], interpretability [Gilpin et al. 2018], fairness [Mehrabi et al. 2021], and predictability [Bansal et al. 2019]. Our general model abstracts away from specific application domains and interfaces, but as with the study of Human-AI collaboration it is crucial that a human user understand the AI's communication protocol and preferences in order to correctly and usefully interpret any given AI interaction.

Information Design and Decision-Making under Imperfect Information. We model AI agents as virtual consultants that provide information to human users. This is

closely related to the impact of changing the amount of information available to a human agent when taking actions or making decisions. For instance, there is a growing line of literature on how to use historical sales data to set prices and/or design the rules of a marketplace or auction [Den Boer 2015; Cole and Roughgarden 2014; Morgenstern and Roughgarden 2016], how buyer and/or seller behavior is impacted by the quantity or quality of data they have access to [Bergemann et al. 2015; Bergemann and Morris 2019], and how mechanism design can be augmented by imperfect predictions [Agrawal et al. 2022; Xu and Lu 2022]. Our model includes such scenarios by assuming that the data is provided by and accessed via an AI agent. The primary difference is that we view an AI agent as having some degree of agency and preference, so that rather than simply passing unfiltered data the AI agent may provide some additional processing or analysis on the user’s behalf.

The notion of filtering data evokes the study of information design [Bergemann and Morris 2016], in which the “designer” strategically selects which information to provide to players in a game. In Bayesian Persuasion, for example, a sender commits to a protocol that determines which (partial) signals to reveal to a receiver, in order to induce a desirable behavior [Kamenica and Gentzkow 2011; Kamenica 2019]. Our model differs in that the AI agent does not have a stake in the action ultimately taken by its human user, so there can be no temptation to persuade. However, the AI agent may still have (possibly misaligned) incentives over the communication transcript and this can influence the nature of information that is shared. Analyzing the outcome of communication with AI agents therefore shares much in common with analyzing the space of behaviors that can be obtained under different information structures.

Simulating Strategic Agents with AI. There is a recent and rapidly growing line of research using large language models to explicitly simulate the behavior of strategic agents. This includes simulating the outcome of game-theoretic studies in economics [Horton 2023; Shapira et al. 2024], generating virtual panels of AI-powered individuals for market research [Brand et al. 2023], or designing AI agents to play strategic games such as Werewolf or Diplomacy [Xu et al. 2023; Bakhtin et al. 2022]. These studies suggest that modern LLMs have a (perhaps implicit) capacity for strategic reasoning, encoded within their training data, that can be measured and evaluated [Raman et al. 2024]. In a similar vein, a line of work on “common sense reasoning” by AI systems has demonstrated consistent improvement by LLMs on challenge datasets that are designed to encode and test “knowledge that is commonly assumed in other humans” [Davis and Marcus 2015; Sakaguchi et al. 2021; Li et al. 2023; Touvron et al. 2023]. Again, this suggests a notion of common-sense reasoning embedded in LLMs that could be distilled and recovered.

Together, these lines of work support the idea that AI-powered agents can have latent knowledge and behaviors that resemble human economic agents. But we do not restrict our attention to scenarios where an AI system is explicitly directed to simulate a human or an economic agent. Rather, a modern AI implicitly acts as an economic agent for *all* tasks, including tasks that are not directly economic or strategic. The fact that generative AI can effectively simulate human-like common sense and strategic reasoning supports our choice to model them as economic agents.

Designing and Evaluating Generative AI. Underlying the economic analyses of generative AI tools, there is a huge and rapidly growing literature on generative AI itself, and particularly LLMs. The scope is very broad: from designing the AIs (including foundation models and training thereof [Vaswani et al. 2017; Brown et al. 2020; OpenAI 2023; Touvron et al. 2023; Gunasekar et al. 2023] and fine-tuning [Ouyang et al. 2022; Bai et al. 2022]; see [Zhao et al. 2023] for a survey) to orchestrating their usage (e.g., prompting techniques [Brown et al. 2020; Wei et al. 2022; Kojima et al. 2022], integration with tools [Wolfram 2023; Gao et al. 2023]) to evaluating their capabilities [Bubeck et al. 2023; Chang et al. 2024]. While economic models are abstract by nature and hide most implementation details, the models (and the questions being asked) should be grounded in the state-of-art in machine learning: which capabilities are feasible now, what is projected to be realistic (or not) in the near future, and what are the salient costs and tradeoffs. This is particularly important when economic models need to incorporate some domain-specific structures, such as costs or constraints, that may be first-order concerns in a given application scenario.

REFERENCES

- AGRAWAL, P., BALKANSKI, E., GKATZELIS, V., OU, T., AND TAN, X. 2022. Learning-augmented mechanism design: Leveraging predictions for facility location. In *Proceedings of the 23rd ACM Conference on Economics and Computation*. 497–528.
- AIRBNB. 2023. Smart pricing. <https://www.airbnb.com/help/article/1168>.
- AMAZON. 2023. Adjust pricing quickly and automatically. <https://sell.amazon.com/tools/automate-pricing>.
- AMERSHI, S., WELD, D., VORVOREANU, M., FOURNEY, A., NUSHI, B., COLLISSON, P., SUH, J., IQBAL, S., BENNETT, P. N., INKPEN, K., ET AL. 2019. Guidelines for human-ai interaction. In *Proceedings of the 2019 chi conference on human factors in computing systems*. 1–13.
- ARMSTRONG, M. AND VICKERS, J. 2010. A model of delegated project choice. *Econometrica* 78, 1, 213–244.
- BAI, Y., JONES, A., NDOUSSE, K., ASKELL, A., CHEN, A., DASSARMA, N., DRAIN, D., FORT, S., GANGULI, D., HENIGHAN, T., AND ET AL. 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*.
- BAKHTIN, A., BROWN, N., DINAN, E., FARINA, G., FLAHERTY, C., FRIED, D., GOFF, A., GRAY, J., HU, H., (FAIR), M. F. A. R. D. T., ET AL. 2022. Human-level play in the game of diplomacy by combining language models with strategic reasoning. *Science* 378, 6624, 1067–1074.
- BANSAL, G., NUSHI, B., KAMAR, E., WELD, D. S., LASECKI, W. S., AND HORVITZ, E. 2019. Updates in human-ai teams: Understanding and addressing the performance/compatibility tradeoff. In *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 33. 2429–2437.
- BENDOR, J., GLAZER, A., AND HAMMOND, T. 2001. Theories of delegation. *Annual review of political science* 4, 1, 235–269.
- BERGEMANN, D., BROOKS, B., AND MORRIS, S. 2015. The limits of price discrimination. *American Economic Review* 105, 3, 921–957.
- BERGEMANN, D. AND MORRIS, S. 2016. Information design, bayesian persuasion, and bayes correlated equilibrium. *American Economic Review* 106, 5, 586–591.
- BERGEMANN, D. AND MORRIS, S. 2019. Information design: A unified perspective. *Journal of Economic Literature* 57, 1, 44–95.
- BRAND, J., ISRAELI, A., AND NGWE, D. 2023. Using gpt for market research. *Available at SSRN* 4395751.
- BROWN, T., MANN, B., RYDER, N., SUBBIAH, M., KAPLAN, J. D., DHARIWAL, P., NEELAKANTAN, A., SHYAM, P., SASTRY, G., ASKELL, A., AGARWAL, S., HERBERT-VOSS, A., KRUEGER, G., HENIGHAN, T., CHILD, R., RAMESH, A., ZIEGLER, D., WU, J., WINTER, C., HESSE, C., CHEN,

- M., SIGLER, E., LITWIN, M., GRAY, S., CHESS, B., CLARK, J., BERNER, C., MCCANDLISH, S., RADFORD, A., SUTSKEVER, I., AND AMODEI, D. 2020. Language models are few-shot learners. In *34th Advances in Neural Information Processing Systems (NeurIPS)*. Vol. 33. 1877–1901.
- BUBECK, S., CHANDRASEKARAN, V., ELDAN, R., GEHRKE, J., HORVITZ, E., KAMAR, E., LEE, P., LEE, Y. T., LI, Y., LUNDBERG, S., ET AL. 2023. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*.
- CHANG, Y., WANG, X., WANG, J., WU, Y., YANG, L., ZHU, K., CHEN, H., YI, X., WANG, C., WANG, Y., YE, W., ZHANG, Y., CHANG, Y., YU, P. S., YANG, Q., AND XIE, X. 2024. A survey on evaluation of large language models. *ACM Trans. on Intelligent Systems and Technology* 15, 3.
- COLE, R. AND ROUGHGARDEN, T. 2014. The sample complexity of revenue maximization. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*. 243–252.
- DAVIS, E. AND MARCUS, G. 2015. Commonsense reasoning and commonsense knowledge in artificial intelligence. *Communications of the ACM* 58, 9, 92–103.
- DEN BOER, A. V. 2015. Dynamic pricing and learning: historical origins, current research, and new directions. *Surveys in operations research and management science* 20, 1, 1–18.
- DUETTING, P., MIRROKNI, V., LEME, R. P., XU, H., AND ZUO, S. 2023. Mechanism design for large language models. *arXiv preprint arXiv:2310.10826*.
- FISH, S., GONCZAROWSKI, Y. A., AND SHORRER, R. I. 2024. Algorithmic collusion by large language models. *arXiv preprint arXiv:2404.00806*.
- GAO, L., MADAAN, A., ZHOU, S., ALON, U., LIU, P., YANG, Y., CALLAN, J., AND NEUBIG, G. 2023. Pal: Program-aided language models. In *40th Intl. Conf. on Machine Learning (ICML)*. 10764–10799.
- GILPIN, L. H., BAU, D., YUAN, B. Z., BAJWA, A., SPECTER, M., AND KAGAL, L. 2018. Explaining explanations: An overview of interpretability of machine learning. In *2018 IEEE 5th International Conference on data science and advanced analytics (DSAA)*. IEEE, 80–89.
- GOOGLE. 2023. About automated bidding. <https://support.google.com/google-ads/answer/2979071?hl=en>.
- GROSZ, B. J. 1996. Collaborative systems (aaai-94 presidential address). *AI magazine* 17, 2, 67–67.
- GUNASEKAR, S., ZHANG, Y., ANEJA, J., MENDES, C. C. T., GIORNO, A. D., GOPI, S., JAVAHERIPI, M., KAUFFMANN, P., DE ROSA, G., SAARIKIVI, O., SALIM, A., SHAH, S., BEHL, H. S., WANG, X., BUBECK, S., ELDAN, R., KALAI, A. T., LEE, Y. T., AND LI, Y. 2023. Textbooks are all you need. *arXiv preprint arXiv:2306.11644*.
- HOFF, K. A. AND BASHIR, M. 2015. Trust in automation: Integrating empirical evidence on factors that influence trust. *Human factors* 57, 3, 407–434.
- HOLMSTROM, B. 1984. *Bayesian Models in Economic Theory*. Elsevier, Amsterdam, Chapter On the Theory of Delegation.
- HORTON, J. J. 2023. Large language models as simulated economic agents: What can we learn from homo silicus? Tech. rep., National Bureau of Economic Research.
- KAMENICA, E. 2019. Bayesian persuasion and information design. *Annual Review of Economics* 11, 1, 249–272.
- KAMENICA, E. AND GENTZKOW, M. 2011. Bayesian persuasion. *American Economic Review* 101, 6, 2590–2615.
- KOJIMA, T., GU, S. S., REID, M., MATSUO, Y., AND IWASAWA, Y. 2022. Large language models are zero-shot reasoners. *35th Advances in Neural Information Processing Systems (NeurIPS)*.
- LAFFONT, J.-J. AND MARTIMORT, D. 2002. *The Theory of Incentives: The Principal-Agent Model*. Princeton University Press.
- LI, Y., BUBECK, S., ELDAN, R., DEL GIORNO, A., GUNASEKAR, S., AND LEE, Y. T. 2023. Textbooks are all you need ii: phi-1.5 technical report. *arXiv preprint arXiv:2309.05463*.
- MAS-COLELL, A., WHINSTON, M. D., GREEN, J. R., ET AL. 1995. *Microeconomic theory*. Vol. 1. Oxford university press New York.
- MEHRABI, N., MORSTATTER, F., SAXENA, N., LERMAN, K., AND GALSTYAN, A. 2021. A survey on bias and fairness in machine learning. *ACM computing surveys (CSUR)* 54, 6, 1–35.

- META. 2023. About meta bid strategies. <https://www.facebook.com/business/help/1619591734742116?id=2196356200683573>.
- MICROSOFT. 2023. Automated bidding. <https://about.ads.microsoft.com/en-us/solutions/tools/automated-bidding>.
- MORGENSTERN, J. AND ROUGHGARDEN, T. 2016. Learning simple auctions. In *Conference on Learning Theory*. PMLR, 1298–1318.
- O'HARA, M. 2015. High frequency market microstructure. *Journal of financial economics* 116, 2, 257–270.
- OPENAI. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- OUYANG, L., WU, J., JIANG, X., ALMEIDA, D., WAINWRIGHT, C. L., MISHKIN, P., ZHANG, C., AGARWAL, S., SLAMA, K., RAY, A., AND ET AL. 2022. Training language models to follow instructions with human feedback. In *35th Advances in Neural Information Processing Systems (NeurIPS)*.
- RAMAN, N., LUNDY, T., AMOYAL, S., LEVINE, Y., LEYTON-BROWN, K., AND TENNENHOLTZ, M. 2024. Rationality report cards: Assessing the economic rationality of large language models. *arXiv preprint arXiv:2402.09552*.
- SAKAGUCHI, K., BRAS, R. L., BHAGAVATULA, C., AND CHOI, Y. 2021. Winogrande: An adversarial winograd schema challenge at scale. *Communications of the ACM* 64, 9, 99–106.
- SHAPIRA, E., MADMON, O., REICHART, R., AND TENNENHOLTZ, M. 2024. Can large language models replace economic choice prediction labs? *arXiv preprint arXiv:2401.17435*.
- TOUVRON, H., LAVRIL, T., IZACARD, G., MARTINET, X., LACHAUX, M.-A., LACROIX, T., ROZIÈRE, B., GOYAL, N., HAMBRO, E., AZHAR, F., ET AL. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- TOUVRON, H., MARTIN, L., STONE, K., ALBERT, P., ALMAHAIRI, A., BABAEI, Y., BASHLYKOV, N., BATRA, S., BHARGAVA, P., BHOSALE, S., ET AL. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- VASWANI, A., SHAZEER, N., PARMAR, N., USZKOREIT, J., JONES, L., GOMEZ, A. N., KAISER, L., AND POLOSUKHIN, I. 2017. Attention is all you need. In *30th Advances in Neural Information Processing Systems (NeurIPS)*. 5998–6008.
- WEI, J., WANG, X., SCHUURMANS, D., BOSMA, M., XIA, F., CHI, E., LE, Q. V., ZHOU, D., ET AL. 2022. Chain-of-thought prompting elicits reasoning in large language models. *36th Advances in Neural Information Processing Systems (NeurIPS)* 35, 24824–24837.
- WOLFRAM, S. 2023. Chatgpt gets its “wolfram superpowers”! A blog post: <https://writings.stephenwolfram.com/2023/03/chatgpt-gets-its-wolfram-superpowers/>.
- XU, C. AND LU, P. 2022. Mechanism design with predictions. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence*. 571–577.
- XU, Y., WANG, S., LI, P., LUO, F., WANG, X., LIU, W., AND LIU, Y. 2023. Exploring large language models for communication games: An empirical study on werewolf. *arXiv preprint arXiv:2309.04658*.
- YAO, F., LI, C., NEKIPELOV, D., WANG, H., AND XU, H. 2024. Human vs. generative ai in content creation competition: Symbiosis or conflict? *arXiv preprint arXiv:2402.15467*.
- ZHAO, W. X., ZHOU, K., LI, J., TANG, T., WANG, X., HOU, Y., MIN, Y., ZHANG, B., ZHANG, J., DONG, Z., DU, Y., YANG, C., CHEN, Y., CHEN, Z., JIANG, J., REN, R., LI, Y., TANG, X., LIU, Z., LIU, P., NIE, J.-Y., AND WEN, J.-R. 2023. A survey of large language models. *arXiv preprint arXiv:2303.18223*.

A. DETAILED EXAMPLES

A.1 Warm-up: Single-Player Email Game

The following simple scenario is intended as a warm-up and showcases how our framework incorporates delegation. A single human player is writing an email

and is debating between two different styles, A or B . The state of nature $\omega \in \{A, B\}$ determines which email style is most appropriate for the given scenario. The human's prior is that the state is equally likely to be A or B . The human can choose whether to think carefully about the situation, which comes at a cost of 4 but reveals the state of nature. They then choose whether to write email A or B . Choosing the correct style generates value 5 while choosing the incorrect style generates value -10 .

Imagine now that the human player has access to an AI agent. The AI agent has access to a signal $\sigma \in \{A, B\}$ that is equal to the true state with probability 90%.¹² The human player can choose whether to request a signal $m \in \{A, B\}$ from the AI agent. Doing so incurs a cost of 1 for the user. The AI obtains utility 1 if it correctly reports the state of nature and 0 otherwise. Given this utility and information structure, the AI will maximize utility by reporting $m = \sigma$. In this case, the human maximizes utility by requesting a signal m from the AI, choosing not to think carefully, then writing an email in style m . This results in an expected payoff of $(0.9)(5) + (0.1)(-10) - 1 = 2.5$ for the user. Notably, this results in a reduced probability of selecting the appropriate email style.

A.2 Costly Signaling

A professor is writing a recommendation letter for a student. The student is equally likely to be either Typical (T) or Strong (S), and this is encoded in the state of nature $\omega \in \{T, S\}$. There are two human players, a sender (the professor) and a receiver. The sender observes the student's strength, but the receiver does not. The sender first chooses whether to write a weak or strong letter. Writing a weak letter is free but writing a strong letter costs the sender 4. The receiver then observes the sender's letter and chooses whether to hire the student. The receiver obtains payoff 1 for hiring a strong student, payoff -2 for hiring a typical student, and payoff 0 for not hiring. The sender obtains value 10 if a strong student is hired, value 6 if a typical student is hired, and value 0 if the student is not hired.¹³

This game has two equilibria. First, there is a trivial "babbling equilibrium" where the receiver never hires (i.e., ignores the letter) and the sender always writes a weak letter. But there is also a non-trivial equilibrium, which we will describe first in terms of the receiver's strategy and then the sender's strategy. The receiver does not hire upon receiving a weak letter, and hires with probability p upon receiving a strong letter. The sender always writes a strong letter for strong candidates, and writes a strong letter for a typical candidate with probability q . We have an equilibrium at $(p, q) = (2/3, 1/2)$. Indeed, if $q = 1/2$ then, conditional on receiving a strong letter, the receiver's posterior belief is that the student is strong with probability $2/3$, which makes her indifferent between hiring and not hiring. If $p = 2/3$ then the sender strictly prefers to write a strong letter for strong students (expected total payoff $10 \times 2/3 - 4$ for a strong letter versus payoff 0 for a weak letter), and is indifferent between writing strong and typical letters for typical students (expected

¹²Formally, the state of nature is now a pair $(\omega, \sigma) \in \{A, B\}^2$, initially drawn from a joint distribution.

¹³One could obtain a qualitatively similar outcome if the sender's payoff if the student is hired is independent of quality, but writing a strong letter is easier/cheaper for strong students.

total payoff $6 \times \frac{2}{3} - 4 = 0$ for a strong letter). Under this equilibrium, a strong student is hired with probability $\frac{2}{3}$ and a typical student is hired with probability $\frac{1}{2} \times \frac{2}{3} = \frac{1}{3}$. The sender's expected utility is $\frac{1}{2} (10 \times \frac{2}{3} - 4) = \frac{4}{3}$, the surplus obtained from writing strong letters for strong students.

Next imagine that the sender has an AI agent that can assist with the letter-writing process, so that the cost of writing a strong letter effectively reduces from 4 to 1. (The receiver's AI agent, if any, does not provide any assistance in this matter.) If we repeat the equilibrium calculation above, our choice of $q = \frac{1}{2}$ will be unchanged, but making the sender indifferent between strong and weak letters for typical students requires that the receiver considers the letter with probability only $p = \frac{1}{6}$. Thus, at equilibrium, a strong student is hired with probability $\frac{1}{6}$ and a typical student is hired with probability $\frac{1}{2} \times \frac{1}{6} = \frac{1}{12}$. In this equilibrium the sender's utility falls to $\frac{1}{2} (10 \times \frac{1}{6} - 1) = \frac{1}{3}$, which is again the surplus obtained from writing strong letters for strong students. Notably, not only are students hired less often, but the sender's expected utility falls at equilibrium due to the presence of the AI agent. This is because the sender cannot commit to *not* using the AI agent, and this limits the sender's ability to signal the student's type via costly effort.

A.3 Advice Evaluated on Perceived Helpfulness: Numerical Example

The following is an instance of the AI advice game described in Section 2.1. A user is deciding which of two products to purchase and can consult an AI agent. Product A is a safe choice that always results in utility 0 for the user. Product B has a two binary features, $\omega = (\omega_1, \omega_2) \in \{0, 1\}^2$, that are drawn uniformly at random. The realization of ω is known to the AI agent but not the human user. The user has a utility function $v: \{0, 1\}^2 \rightarrow \mathbb{R}$ that maps realizations of ω into a (possibly negative) payoff for purchasing Product B. The user can communicate with the AI agent before making a purchasing decision. Specifically, if the user requests advice, the AI sends a message that contains one (and only one) of the product features.¹⁴

Suppose that, motivated by a design objective of maximizing perceived helpfulness, the AI agent's payoffs are determined as follows. Declaring a feature incorrectly gives payoff $-\infty$ for the AI. Otherwise, the AI agent's payoff is the difference in expected utility between the user's optimal choice given the revealed feature(s) and their optimal choice given no information, where the expectation is over any features that are not revealed. This payoff is always non-negative, and is 0 whenever the AI agent's advice would not influence a utility-maximizing user's choice of action.

Suppose that the user's utility function is $v(1, 1) = 2$ and $v(0, 0) = v(0, 1) = v(1, 0) = -1$. I.e., the user gets utility 2 for buying B if both features are 1, otherwise utility -1 . In this case, absent any information from the AI agent, the expected utility-maximizing action is to purchase A since $\frac{3}{4} \cdot (-1) + \frac{1}{4} \cdot 2 < 0$. If either feature is revealed to be 0, then the user's utility is certainly -1 for purchasing B so the optimal choice is still purchase A. But if either feature is revealed to be

¹⁴The restriction to declaring only one feature is motivated by communication length; in general an AI agent cannot practically communicate all information it has access to.

1, the expected utility from purchasing B (in expectation over the other feature) is $\frac{1}{2} \cdot (-1) + \frac{1}{2} \cdot 2 = \frac{1}{2} > 0$. This means that correctly revealing a feature to be 1 yields payoff $\frac{1}{2}$ for the AI, whereas revealing a feature to be 0 yields payoff 0. We emphasize that this payoff to the AI is independent of which product the user ultimately decides to purchase; it is a function only of ω and the communication transcript.

This situation creates misaligned incentives for the human user and the AI. If the realized state is $\omega = (1, 0)$, then the user's ex post preferred action is to purchase A. But under this realization the AI is incentivized to reveal $\omega_1 = 1$. In fact, the AI is incentivized to reveal a feature with value 1 in all states except $\omega = (0, 0)$. What this means is that, conditional on the AI revealing a feature $\omega_i = 1$, a rational user should infer only that $\omega \neq (0, 0)$, so her expected utility from purchasing B is $\frac{2}{3} \cdot (-1) + \frac{1}{3} \cdot 2 = 0$. In other words, the AI agent confers zero value even to a fully rational user that correctly anticipates its behavior.

A.4 Delegated Search: Numerical Example

The following is a toy example of the AI delegated search game described in Section 2.2. The example illustrates a scenario where (a) a user can potentially benefit by strategically inflating the strength of their preferences in order to bypass the AI agent's preference to avoid certain outcome, but (b) this can lead to suboptimal outcomes when the set of possible outcomes available to the AI agent is not perfectly understood by the user.

In this toy example, a user is writing an email with the assistance of an AI agent. Emails are evaluated on two dimensions, say humor and professionalism. Under this parameterization, $\Gamma = [0, 1]^2$ denotes the space of all conceivable emails, where $(1, 0)$ is a perfectly humorous but unprofessional email and $(0, 1)$ is perfectly professional but not at all funny. Not all points in the parameter space are necessarily achievable, so while $(1, 1)$ represents a hypothetical email that is both perfectly funny and perfectly professional, this may not be possible to achieve. Here $X \subseteq [0, 1]^2$ is the set of all achievable parameters.

The user's utility from an email is characterized by a relative preference for humor versus professionalism. This is described by a vector $y \in \mathbb{R}_{\geq 0}^2$ with $\|y\|_2 = 1$. An email $x \in \Gamma$ then generates utility $u_y(x) = x \cdot y$, where we write u_y for the utility function corresponding to preference vector y . We let U be the space of all such utility functions.

Let's say the user knows that X contains $x^f = (1, 0)$ and $x^p = (0, 1)$ for "funny" and "professional", respectively, so emails achieving these parameters exist. Let's further imagine that the AI agent is designed to downplay humor unless specifically requested; this is encoded as a penalty function with $\gamma(x^f) = 0.2$ and $\gamma(x^p) = 0$. One can then verify that if $X = \{x^f, x^p\}$, the AI agent will return x^p in response to declared relative preferences $y = (y_1, y_2)$ if $y_1 < 0.8$ (and hence $y_2 > 0.6$), and will return x^f if $y_1 > 0.8$. Thus, if the user has a strict preference for the humorous email (no matter how slight), they will optimize by declaring a strong preference for humor.

Next imagine that the AI agent is also able to create emails that combine professionalism and humor to a certain extent. Concretely, there are points $x^m = (2/3, 2/3)$ ('m' for mediocre at both) and $x^{pf} = (1/2, 1)$ ('pf' for professional yet funny) in the

set X as well. But the AI agent steers away from including humor in professional emails; say this is encoded with penalties $\gamma(x^m) = 0.25$ and $\gamma(x^{pf}) = 0.35$. These numbers are chosen so that the threshold on the declared preference for humor at which x^m is returned, versus x^{pf} , is $\approx 0.74 < 0.8$.

A feature of this example is that if the agent's most-preferred option is x^{pf} ,¹⁵ then it would be optimal to declare a moderate preference for humor; say $y_1 = 0.71$. On the other hand, if the user is unaware that this option is possible and prefers x^f to x^p ,¹⁶ they would rationally declare a stronger preference for humor, $y_1 > 0.8$, which would be suboptimal with respect to the true choice set. This toy example demonstrates that while an agent can benefit by strategically declaring their preferences to an AI that attempts to steer away from certain outcomes, this can lead to suboptimal outcomes when the set of possible options is not perfectly understood. While this example is described in terms of an incorrect belief about X , the same conclusion can be reached if the user has a prior over X , in which case a suboptimal choice is reached on some realizations of X .

¹⁵For example, this is the case if the agent's true preference is given by $z = (z_1, z_2)$ with $z_1 = 0.79$.

¹⁶Which is likewise obtained by true preferences $z = (z_1, z_2)$ where $z_1 = 0.79$

Causal Inference under Incentives: An Annotated Reading List

KEEGAN HARRIS

Carnegie Mellon University

and

VASILIS SYRGKANIS

Stanford University

We provide an overview of research on causal inference in the presence of strategic agents. Work in this area uses tools from econometrics, statistics, machine learning, and game theory to infer causal relationships between treatments and outcomes of interest when the treated individuals have an incentive to behave strategically.

Categories and Subject Descriptors: F.7.2 [**Theory of computation**]: Algorithmic game theory; K.4 [**Computing methodologies**]: Machine learning

General Terms: Causal Inference, Game Theory, Machine Learning

Additional Key Words and Phrases: Decision-making, Incentives

Learning causal relationships from data is an important task across a wide variety of domains ranging from healthcare and drug development, to online advertising and e-commerce. As a result, there has been much work in the literature on economics, statistics, computer science, and public policy on designing algorithms and methodologies for causal inference.

While most of the focus has been on questions which are statistical in nature, one must also take game-theoretic incentives into consideration when doing causal inference about strategic individuals who have a preference over the treatment they receive. For example, it may be hard to infer causal relationships in randomized control trials when there is non-compliance by participants in the study. More generally, causal learning may be difficult whenever individuals are free to self-select their own treatments and there is sufficient heterogeneity between individuals with different preferences. Even when compliance can be enforced, individuals may strategize by modifying the attributes they present to the causal inference process in order to be assigned a more desirable treatment.

This annotated reading list is intended to serve as a brief summary of work on causal inference in the presence of strategic agents. While this list is not comprehensive, we hope that it will be a useful starting point for members of the SIGecom community to learn more about this exciting research area at the intersection of causal inference, game theory, and machine learning.

The reading list is organized as follows: (1, 3) study non-compliance in randomized trials, (2-4) focus on instrumental variable methods, (4-6) consider incentive misalignment between the individual running the causal inference procedure and

Authors' addresses: keeganh@cmu.edu, vsyrgk@stanford.edu

the subjects of the procedure, (7,8) study cross-unit interference, and (9,10) are about synthetic control methods.

- (1) [Robins 1998]: This paper provides an overview of methods to correct for non-compliance in randomized trials (i.e., non-adherence by trial participants to the treatment assignment protocol).
- (2) [Angrist et al. 1996]: This seminal paper outlines the concept of instrumental variables (IVs) and describes how they can be used to estimate causal effects. An IV is a variable that affects the treatment variable but is unrelated to the outcome variable except through its effect on the treatment. IV methods leverage the fact that variation in IVs is independent of any confounding to estimate the causal effect of the treatment.
- (3) [Ngo et al. 2021]: Unlike prior work on non-compliance in clinical trials, this work leverages tools from information design to reveal information about the effectiveness of the treatments in such a way that participants become incentivized to comply with the treatment recommendations over time.
- (4) [Harris et al. 2022]: This paper studies the problem of making decisions about a population of strategic agents. The authors make the novel observation that the assessment rule deployed by the principal is a valid instrument, which allows them to apply standard methods for instrumental variable regression to learn causal relationships in the presence of strategic behavior.
- (5) [Miller et al. 2020]: This paper considers the problem of strategic classification, where a principal makes decisions about a population of strategic agents. Given knowledge of the principal’s deployed assessment rule, the agents may strategically modify their observable features in order to receive a more desirable assessment. The authors are the first to show that designing good incentives for agent improvement (i.e. encouraging strategizing in a way which actually benefits the agent) is at least as hard as orienting edges in the corresponding causal graph.
- (6) [Wang et al. 2023]: Incentive misalignment between patients and providers may occur when average treated outcomes are used as quality metrics. Such misalignment is generally undesirable in healthcare domains, as it may lead to decreased patient welfare. To mitigate this issue, this work proposes an alternative quality metric, the total treatment effect, which accounts for counterfactual untreated outcomes. The authors show that rewarding the total treatment effect maximizes total patient welfare.
- (7) [Wager and Xu 2021]: Motivated by applications such as ride-sharing and tuition subsidies, this work studies settings in which interventions on one unit may have effects on others (i.e., cross-unit interference). The authors focus on the problem of setting supply-side payments in a centralized marketplace. They use a mean-field modeling-based approach to model the cross-unit interference, and design a class of experimentation schemes which allow them to optimize payments without disturbing the market equilibrium.
- (8) [Li et al. 2023]: Like [Wager and Xu 2021], this paper studies the effects of cross-unit interference, although the interference considered here comes from

congestion in a service system. As a result, the interference considered here is dynamic, in contrast to the static interference considered in the previous entry.

- (9) [Abadie and Gardeazabal 2003]: This is the first paper on synthetic control methods (SCMs), a popular technique for estimating counterfactuals from panel data. In the SCM setup, there is a pre-intervention time period during which all units are under control, followed by a post-intervention time period when all units undergo exactly one intervention (either the treatment or control). Given a test unit (who was given the treatment) and a set of donor units (who remained under control), SCMs use the pre-treatment data to learn a relationship (usually linear or convex) between the test and donor units. This relationship is then extrapolated to the post-intervention time period in order to estimate the counterfactual trajectory for the test unit under control.
- (10) [Ngo et al. 2023]: A common assumption in the literature on SCMs is that of “overlap”: the outcomes for the test unit can be written as a combination (e.g., linear or convex) of the donor units. This work sheds light on this often overlooked assumption and shows that (i) when units select their own treatments and (ii) there is sufficient heterogeneity between units who prefer different treatments, then overlap does not hold. Like [Ngo et al. 2021], the authors use tools from information design and multi-armed bandits to incentivize units to explore different treatments in a way which ensures that the overlap condition will gradually become satisfied over time.

REFERENCES

- ABADIE, A. AND GARDEAZABAL, J. 2003. The economic costs of conflict: A case study of the basque country. *American economic review* 93, 1, 113–132.
- ANGRIST, J. D., IMBENS, G. W., AND RUBIN, D. B. 1996. Identification of causal effects using instrumental variables. *Journal of the American statistical Association* 91, 434, 444–455.
- HARRIS, K., NGO, D. D. T., STAPLETON, L., HEIDARI, H., AND WU, S. 2022. Strategic instrumental variable regression: Recovering causal relationships from strategic responses. In *International Conference on Machine Learning*. PMLR, 8502–8522.
- LI, S., JOHARI, R., KUANG, X., AND WAGER, S. 2023. Experimenting under stochastic congestion. *arXiv preprint arXiv:2302.12093*.
- MILLER, J., MILLI, S., AND HARDT, M. 2020. Strategic classification is causal modeling in disguise. In *International Conference on Machine Learning*. PMLR, 6917–6926.
- NGO, D., HARRIS, K., AGARWAL, A., SYRGKANIS, V., AND WU, Z. S. 2023. Incentive-aware synthetic control: Accurate counterfactual estimation via incentivized exploration. *arXiv preprint arXiv:2312.16307*.
- NGO, D. D. T., STAPLETON, L., SYRGKANIS, V., AND WU, S. 2021. Incentivizing compliance with algorithmic instruments. In *International Conference on Machine Learning*. PMLR, 8045–8055.
- ROBINS, J. M. 1998. Correction for non-compliance in equivalence trials. *Statistics in medicine* 17, 3, 269–302.
- WAGER, S. AND XU, K. 2021. Experimenting in equilibrium. *Management Science* 67, 11, 6694–6715.
- WANG, S., BATES, S., ARONOW, P., AND JORDAN, M. I. 2023. Operationalizing counterfactual metrics: Incentives, ranking, and information asymmetry. *arXiv preprint arXiv:2305.14595*.

Impartial Peer Selection: An Annotated Reading List

OMER LEV

Ben-Gurion University

and

HARPER LYON

Tulane University

and

NICHOLAS MATTEI

Tulane University

The study of peer selection mechanisms presents a unique opportunity to understand and improve the practice of a group selecting its best members, despite each member of that group wanting to be selected. A prime example of such a setting is academic peer review, for which peer selection offers a variety of improvement directions. We present an annotated reading list covering the foundations of peer selection as well as recent and emerging work within the field.

Categories and Subject Descriptors: B.6.3 [**Logic Design**]: Theory of computing—*Algorithmic mechanism design*

General Terms: Algorithms, Performance, Theory

Additional Key Words and Phrases: Computational social choice, mechanism design, strategic agents

Peer selection refers to any social choice problem where agents are asked to select some subset of themselves to receive an award or benefit, which each of them would like to receive. Perhaps most notably, this serves as a (worst case) model for the way academic conferences conduct peer review, a key element of modern science.

Academic reviews and grant boards have been the subject of many empirical studies that, among other issues, have focused on the effectiveness and limits of the system, often substantiating anecdotal complaints of bias and inaccuracy [Cole et al. 1981; McNutt et al. 1990; Wenneras and Wold 1997]. These conclusions, as well as demands of practicality, suggest that broadening the base of reviewers is essential for improving the process of peer review. However, increasing the size of the reviewer pool is not without serious challenges, not the least of which is the sourcing of additional reviewers. A common solution – also proposed by the 2009 National Science Foundation’s (NSF) Mechanism Design Proposal Pilot – is to require those that submit proposals to also act as reviewers, casting the problem of peer review firmly into the domain of peer selection.

Moving beyond a relatively small group of impartial experts requires robust mechanisms for soliciting and aggregating peer evaluations. Not only to support larger volumes of reviews but also because involving self-interested agents in the review process introduces new game-theoretic challenges. While we still want to select the best agents/work, we also need to worry about reviewers strategically manipulating

their reviews to increase their own chances of being selected. We call a mechanism which does not incentivize strategic reviews *impartial* or *strategyproof*. The papers below highlight both the core theoretical questions underlying the tension between selecting the best work and impartiality, as well as potential mechanisms for better, more robust peer selection. In addition to the papers described below, Olckers and Walsh [2022] provides a comprehensive overview many peer review algorithms and Shah [2022] is a deeper dive into the practical considerations of conference peer review.

- (1) MERRIFIELD, M. R. AND SAARI, D. G. 2009. Telescope time without tears: a distributed approach to peer review. *Astronomy & Geophysics* 50, 4, 4-16

Motivated by the overwhelming number of applications to powerful telescopes, this paper presents an early application of mechanism design to the problem of peer evaluation, and uses a Borda count across submitters' rankings of a subset of proposals to determine acceptance. The mechanism also rewards reviewers according to the similarity of their reviews to the aggregate ranking, aiming to incentivize high quality reviews. This is not an impartial mechanism [Ardabili and Liu 2013], since agents are rewarded for reviews that agree with the community consensus, but the NSF using it as a basis for a pilot program, spurred research in this domain.

- (2) ALON, N., FISCHER, F. A., PROCACCIA, A. D., AND TENNENHOLTZ, M. 2011. Sum of us: strategyproof selection from the selectors. In *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-2011), Groningen, The Netherlands, July 12-14, 2011*, K. R. Apt, Ed. ACM, 101-110

This paper explores the trade-off between impartiality and optimality in peer selection when we have approval votes. Alon et al. prove the impossibility of any deterministic impartial mechanism that finitely approximate the optimal selection, as well as to determine general bounds for randomized impartial mechanisms within the approval voting context. This work initiated a strand of research into partition mechanisms as a means of achieving good approximation guarantees, particularly when selecting a single agent [Holzman and Moulin 2013; Fischer and Klimm 2015; Bousquet et al. 2014]

- (3) KUROKAWA, D., LEV, O., MORGENSTERN, J., AND PROCACCIA, A. D. 2015. Impartial peer review. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015*, Buenos Aires, Argentina, July 25-31, 2015, Q. Yang and M. J. Wooldridge, Eds. AAAI Press, 582-588

With the goal of expanding impartial mechanisms beyond the approval voting context and into settings where reviewers submit rankings or scores for proposals, this paper presents CREDIBLESUBSET. The CREDIBLESUBSET is an impartial peer selection mechanism which removes the potential for manipulation by building a set of all agents who might have been selected had they

strategically manipulated and then randomly selecting agents from this set. This mechanism is proven to be impartial and provides a $k/(k+m)$ approximation of their VANILLA mechanism (which is a randomized version of Borda), where k is the number of desired selections and m is the number of proposals each reviewer is assigned. However, to maintain impartiality, there is a probability that the mechanism returns no selection (or returns some other default option, not based on the evaluation of peers).

- (4) XU, Y., ZHAO, H., SHI, X., AND SHAH, N. B. 2019. On strategyproof conference peer review. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*. IJCAI'19. AAAI Press, 616-622

This paper is interested in creating an impartial reviewer assignment when we have a known set of reviewer conflicts and presents DIVIDE-AND-RANK, which assigns reviewers in a way that avoids conflicts of interests but maintains as much “expertise to review” as possible. As long as the conflict graph obeys certain size and density properties, this graph contraction-based approach is shown to be impartial and is shown to perform well empirically on ICLR submissions and reviews.

- (5) AZIZ, H., LEV, O., MATTEI, N., ROSENSCHEIN, J. S., AND WALSH, T. 2019. Strategyproof peer selection using randomization, partitioning, and apportionment. *Artif. Intell.* 275, 295-309

One shortcoming of partition based mechanisms is their susceptibility to adversarial partitioning, i.e., if all the best proposals appear on one partition. To address this issue this paper presents EXACTDOLLARPARTITION, a randomized impartial mechanism in which agents are partitioned into clusters, and the number of agents selected from each cluster is apportioned by the relative scores of the agents in that cluster. While the agents selected from each cluster are the highest-ranked ones, the key insight is that the total number can vary, avoiding some problems with adversarial partitioning.

- (6) JECMEN, S., ZHANG, H., LIU, R., SHAH, N. B., CONITZER, V., AND FANG, F. 2020. Mitigating manipulation in peer review via randomized reviewer assignments. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds

This paper is concerned with the more concrete problem of minimizing reviewer conflict of interest, in the interpersonal sense, when assigning conference review loads. As a result, the paper focuses on the more practical and empirical aspects of peer review, such as mismatched reviewer expertise, miscalibration between reviewer scores, and the role of group dynamics in shaping discussion. While much of this work is outside of the scope of this reading list, it is an important area of related research. Some additional examples of this style of work can be

found in Ardabili and Liu [2013] and Goldberg et al. [2023].

- (7) NOOTHIGATTU, R., SHAH, N. B., AND PROCACCIA, A. D. 2021. Loss functions, axioms, and peer review. *J. Artif. Intell. Res.* 70, 1481-1515

Multi-category reviews, where reviewers evaluate papers on multiple factors such as significance and readability, are common in real-life peer selection, but aggregating the scores across factors can be challenging. This paper takes a machine learning/learning to rank perspective to reproduce implicit community ranking standards. By using empirical risk minimization and extending the standard L_p norms to multidimensional reviews for use as a loss function, they show that aggregation mechanism derived from the equivalent of the L_1 norm uniquely produce an impartial mechanism with good recall of the original rankings.

- (8) CEMBRANO, J., FISCHER, F. A., HANNON, D., AND KLIMM, M. 2022. Impartial selection with additive guarantees via iterated deletion. In *EC '22: The 23rd ACM Conference on Economics and Computation, Boulder, CO, USA, July 11 – 15, 2022*, D. M. Pennock, I. Segal, and S. Seuken, Eds. ACM, 1104-1105

Taking inspiration from elimination style voting mechanisms like IRV and seeking mechanisms with good approximations bounds for optimal peer selection, this paper describes a “twin threshold” mechanism for approval voting peer selection, where approval votes are removed from agents below a threshold of approval and then selects the agents receiving the maximum non-deleted votes, provided that the agent received more votes than the second, higher, threshold. This mechanism is shown to be impartial and achieves an additive approximation guarantee of $O(n^{1+k}2)$, for n the number of agents and $k \in [0, 1]$, but only when agents’ number of approval votes can be bounded by $O(n^k)$.

- (9) AZIZ, H., MICHA, E., AND SHAH, N. 2023. Group fairness in peer review. In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 – 16, 2023*, A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, Eds

Uniquely for this list, this paper is concerned with a problem other than impartial reviews, as Aziz et al. consider the problem of balancing review assignments for large conferences between research communities to avoid requiring reviewers to evaluate papers far outside of their specialties. Their approach is to extend the game theoretic notion of core fairness to the review assignment setting, and present COBRA, an algorithm which determines a core-respecting assignment on any set of separable and consistent preferences over the assignments.

- (10) LEV, O., MATTEI, N., TURRINI, P., AND ZHYDKOV, S. 2023. Peernomina-
ACM SIGecom Exchanges, Vol. 22, No. 1, June 2024, Pages 113–117

tion: A novel peer selection algorithm to handle strategic and noisy assessments. *Artif. Intell.* 316, 103843

It is an unfortunate reality that reviewers may be imperfectly accurate when it comes to selecting the best papers, but how do we deal with this noise? To address this issue this paper presents PEERNOMINATION, an impartial peer selection mechanism that is able to account for the presence of noisy and/or unreliable reviewers. This mechanism dynamically weighs reviewers' scores, comparing them to other reviewers of the same agents, thus reducing their impact on the outcome. In order to maintain the mechanism's impartiality despite this new manipulation direction, the reviewer assignments have to follow a particular structure, for which an algorithm is also given.

Acknowledgements

Harper Lyon and Nicholas Mattei was supported in part by NSF Awards IIS-RI-2007955, IIS-III-2107505, and IIS-RI-2134857 as well as the Tulane University Jurist Center for Artificial Intelligence and the Tulane University Center for Community-Engaged Artificial Intelligence. Omer Lev was supported, in part, by NSF-BSF grant #2021659, and by Israel Science Fund (ISF) grants #1965/20 and #3152/20.

REFERENCES

- ARDABILI, P. N. AND LIU, M. 2013. Incentives, quality, and risks: A look into the NSF proposal review pilot. *CoRR abs/1307.6528*.
- BOUSQUET, N., NORIN, S., AND VETTA, A. 2014. A near-optimal mechanism for impartial selection. In *Proceedings of the 10th International Conference on Web and Internet Economics (WINE)*. Beijing, China, 133–146.
- COLE, S., COLE, J., AND SIMON, G. 1981. Chance and consensus in peer review. *Science* 214, 4523, 881–886.
- FISCHER, F. AND KLIMM, M. 2015. Optimal impartial selection. *SIAM Journal on Computing* 44, 5, 1263–1285.
- GOLDBERG, A., STELMAKH, I., CHO, K., OH, A. H., AGARWAL, A., BELGRAVE, D., AND SHAH, N. B. 2023. Peer reviews of peer reviews: A randomized controlled trial and other experiments. *CoRR abs/2311.09497*.
- HOLZMAN, R. AND MOULIN, H. 2013. Impartial nominations for a prize. *Econometrica* 81, 1 (January), 173–196.
- McNUTT, R. A., EVANS, A. T., FLETCHER, R. H., AND FLETCHER, S. W. 1990. The effects of blinding on the quality of peer review: A randomized trial. *The Journal of the American Medical Association* 263, 10, 1371–1376.
- OLCKERS, M. AND WALSH, T. 2022. Manipulation and peer mechanisms: A survey. *CoRR abs/2210.01984*.
- SHAH, N. B. 2022. Challenges, experiments, and computational solutions in peer review. *Commun. ACM* 65, 6, 76–87.
- WENNERAS, C. AND WOLD, A. 1997. Nepotism and sexism in peer-review. *Nature* 387, 341–343.

Assortment Optimization: An Annotated Reading Assortment

WILL MA

Columbia University

Which varieties or brands of a product should a retailer stock on its shelf? Carrying a large variety caters to more customers' needs, but could cannibalize the sales of high-end brands and also cause an inventory nightmare. Assortment optimization aims to formalize these tradeoffs, with the basic problem being as follows. There is a universe of brands $j \in U$, each with a market-accepted price r_j . For any $S \subseteq U$, a function $\phi(j, S)$ indicates the probability that a representative customer from the population would purchase j when given the choice from *assortment* S , under the market prices. The optimization problem is to maximize the average revenue per customer, i.e.

$$\max_S \sum_{j \in S} r_j \phi(j, S), \quad (1)$$

possibly with constraints on S due to shelf size. Assortment optimization started out by showing how to efficiently find the optimal S from the exponentially many possibilities, under well-established parametric forms for the function ϕ that are called (*discrete*) *choice models*. Since then, the literature has developed choice models of its own that are specialized for assortment optimization. The basic problem has also been extended, and connected with topics such as online algorithms, machine learning, and mechanism design that are mainstream in the Economics and Computation community, with a vast horizon for future directions.

This is an annotated reading list about assortment optimization, that aims to provide broad coverage while facing a “cardinality constraint” on the number of papers in the assortment.

1. INITIAL THEORY ON EXISTING CHOICE MODELS

- (1) [Talluri and Van Ryzin 2004] An early work to maximize revenue on discrete choice models. Although this paper focuses on the online revenue management problem, a key result attributed to it is that even in the basic problem (1), the optimal assortment has a *revenue-ordered* structure when ϕ falls under the *Multi-Nomial Logit (MNL)* choice model. To elaborate, revenue-ordered means that $S = \{j \in U : r_j \geq \tau\}$ for some price threshold τ . Meanwhile, MNL imposes that $\phi(j, S) = w_j / \sum_{j' \in S \cup \{0\}} w_{j'}$ for some popularity weights $w_{j'}$. MNL is the simplest choice model to capture a non-trivial form of *cannibalization*, where inserting j' into the assortment S decreases the probability of the customer choosing any other j , and 0 represents the “no-purchase” option whose weight is often normalized to 1. This paper also discusses how to estimate the popularity weights in MNL.
- (2) [Rusmevichientong et al. 2010] An early paper to consider constraints on the assortment S , motivated by limited shelf size. The authors develop a search technique that solves cardinality-constrained assortment optimization for MNL. Importantly, they show that the optimal assortment may no longer be revenue-ordered if there is a cardinality constraint, contrasting the structural result of

Author's address: wm2428@gsb.columbia.edu

- [Talluri and Van Ryzin 2004]. This paper also studies how to learn the optimal cardinality-constrained assortment when the function ϕ is initially unknown.
- (3) [Gallego and Topaloglu 2014] Nested Logit is a more general parametric choice model that increases the expressiveness of MNL by allowing for products to be categorized into nests, and having more cannibalization within nest than across nests. This paper makes substantial progress on assortment optimization for Nested Logit, developing a Linear Program to solve the problem even with a cardinality constraint on S . It is also one of the early papers to take an approximation algorithms approach to assortment optimization, by developing a 2-approximation under the more general knapsack constraints. Finally, this paper shows how pricing decisions can be captured using assortment optimization, by creating copies of products with different prices r_j and adding constraints that only one price level for each product can be offered.

2. NEW CHOICE MODELS FOR ASSORTMENT OPTIMIZATION

Parametric choice models have nice analytical forms for the function ϕ , that are also explained by customers drawing random valuations $V_{j'}$ for each $j' \in U \cup \{0\}$ and defining $\phi(j, S) = \Pr[j = \arg\max_{j' \in S \cup \{0\}} V_{j'}]$. However, model selection can be difficult—too few parameters and your demand is misspecified; too many parameters and you overfit. These papers propose non-parametric choice models that focus on prescribing assortment decisions, without worrying about how explainable ϕ is from random-utility theory.

- (4) [Farias et al. 2013] This paper proposes a paradigm for choice modeling where customers have a latent distribution of *ordinal* rankings over $U \cup \{0\}$, motivated by the fact that transaction data in practice only involves a customer making comparisons. The authors take a robust approach to estimating the ranking distribution and solving the assortment optimization problem, that automatically tunes model complexity based on the data.
- (5) [Blanchet et al. 2016] Also motivated by the challenge in model selection, this paper proposes to only capture a customer's first two choices in a Markov chain, and shows how this can simultaneously approximate all random-utility discrete choice models. This surprising insight gives birth to the Markov Chain choice model, and the authors also show how to solve assortment optimization on it.

3. ONLINE VARIANTS

These papers consider the dynamics of multiple sales from different assortments.

- (6) [Golrezaei et al. 2014] An assortment is classically interpreted as a deliberate set of products carried by a brick-and-mortar retailer. This paper considers the personalized assortments that can be offered by an online retailer, and shows how to adjust these based on remaining inventories. It introduces primal-dual analysis and competitive ratios to the assortment optimization literature, and initiated a large body of work that incorporates assortment optimization into online algorithms.
- (7) [Goyal et al. 2016] This paper revitalizes the *dynamic substitution model* where one must jointly decide the assortment S and how much of each $j \in S$ to stock,

which was one of the original motivations of assortment optimization stemming from inventory theory. A stochastic sequence of customers arrives, choosing from the subset of j within S that have remaining inventory. In this model the decision is initial inventory after which the assortment cannot be controlled, contrasting the model of [Golrezaei et al. 2014] where initial inventories are given but the decision is how to dynamically control assortments. This paper develops NP-hardness and PTAS type results for a new formulation where the objective is to maximize revenue subject to a constraint on the total number of units initially stocked, which would become the subject of several follow-ups.

- (8) [Agrawal et al. 2019] This paper considers the joint learning and optimization problem facing a stream of customers who choose according to the same unknown MNL model. It introduces a neat analysis for exploring and exploiting at the same time, illustrating the richness that choice models bring to the classical exploration-exploitation tradeoff and spawning a whole literature on “MNL-Bandit”.

4. FUTURE DIRECTIONS

- (9) [Aouad and Désir 2022] Motivated by the success of deep learning in prediction, this paper proposes an architecture that uses neural networks to accurately predict customer choices. It leverages high-dimensional contextual information about customers and products, while preserving the random-utility structure underlying most choice models. Although this paper does not directly optimize assortments, it plants the seeds for a potentially rich literature that relates assortments to the optimization of inputs to trained machine learning models.
- (10) [Ma 2023] Assortment optimization has been fixated on solving problem (1), which can be interpreted as posting a single assortment for the customer to choose from. But what about more general selling methods? This paper captures assortment optimization as part of a more general Bayesian mechanism design problem, in which customers have ordinal preferences over fixed-price items. It characterizes choice models for which the optimal mechanism is or is not a posted assortment. More generally, assortment optimization motivates the design of Bayesian mechanisms without arbitrary payments, which is potentially a large ground for innovation in the Economics and Computation community.

REFERENCES

- AGRAWAL, S., AVADHANULA, V., GOYAL, V., AND ZEEVI, A. 2019. Mnl-bandit: A dynamic learning approach to assortment selection. *Operations Research* 67, 5, 1453–1485.
- AOUAD, A. AND DÉSIR, A. 2022. Representing random utility choice models with neural networks. *arXiv preprint arXiv:2207.12877*.
- BLANCHET, J., GALLEGO, G., AND GOYAL, V. 2016. A markov chain approximation to choice modeling. *Operations Research* 64, 4, 886–905.
- FARIAS, V. F., JAGABATHULA, S., AND SHAH, D. 2013. A nonparametric approach to modeling choice with limited data. *Management science* 59, 2, 305–322.
- GALLEGO, G. AND TOPALOGLU, H. 2014. Constrained assortment optimization for the nested logit model. *Management Science* 60, 10, 2583–2601.
- GOLREZAEI, N., NAZERZADEH, H., AND RUSMEVICHIENTONG, P. 2014. Real-time optimization of personalized assortments. *Management Science* 60, 6, 1532–1551.

- GOYAL, V., LEVI, R., AND SEGEV, D. 2016. Near-optimal algorithms for the assortment planning problem under dynamic substitution and stochastic demand. *Operations Research* 64, 1, 219–235.
- MA, W. 2023. When is assortment optimization optimal? *Management Science* 69, 4, 2088–2105.
- RUSMEVICHIENTONG, P., SHEN, Z.-J. M., AND SHMOYS, D. B. 2010. Dynamic assortment optimization with a multinomial logit choice model and capacity constraint. *Operations research* 58, 6, 1666–1680.
- TALLURI, K. AND VAN RYZIN, G. 2004. Revenue management under a general discrete choice model of consumer behavior. *Management Science* 50, 1, 15–33.

Recent Trends in Information Elicitation

RAFAEL FRONGILLO

University of Colorado Boulder

and

BO WAGGONER

University of Colorado Boulder

This note provides a survey for the Economics and Computation community of some recent trends in the field of information elicitation. At its core, the field concerns the design of incentives for strategic agents to provide accurate and truthful information. Such incentives are formalized as *proper scoring rules*, and turn out to be the same object as *loss functions* in machine-learning settings, providing many connections. More broadly, the field concerns the design of mechanisms to obtain information from groups of agents and aggregate it or use it for decision making. Recently, work on information elicitation has expanded and been connected to online no-regret learning, mechanism design, fair division, and more.

1. BACKGROUND: SCORING RULES AND PROPERTY ELICITATION

1.1 Introduction

This note surveys recent trends in *information elicitation*, the design of incentives for strategic agents to provide information. Before discussing these recent developments, we recall the key tools of the field, proper scoring rules and property elicitation. These can be viewed as *single-agent* incentives; we then recall multi-agent settings including prediction markets, wagering mechanisms, and peer prediction.

1.2 Single-agent elicitation

The original single-agent elicitation problem is the design of *proper scoring rules* [Brier 1950; Good 1952; Savage 1971; Gneiting and Raftery 2007]. An agent predicts a probability distribution p , then an outcome y is observed. A scoring rule S is a function assigning a score $S(p, y)$ to the prediction p on the observation y . S is *proper* if the agent maximizes expected score, over the randomness in y , by predicting their true belief. For example, the *log score* $S(p, y) = \log p(y)$ is proper.¹

Scoring rules have been extended to the case where the agent predicts the mean, mode, variance, or any other *property* $\Gamma(p)$ of the distribution p of y [Osband 1985; Lambert et al. 2008]; see Gneiting [2011] for the origins of these ideas. The score $S(r, y)$ is said to *elicit* the property Γ if, when the agent's true belief is the distribution p , they uniquely maximize expected score by reporting $r = \Gamma(p)$. For example, the quadratic score $S(r, y) = -\|r - y\|_2^2$, where y and r take values in Euclidean space, elicits the *mean*: the score is maximized by setting $r = \mathbb{E}_p Y$. As

¹To see that the log score is proper, note that the regret of reporting q instead of the true p is $\mathbb{E}_p S(p, Y) - \mathbb{E}_p S(q, Y) = \text{KL}(p \| q)$, which is nonnegative and uniquely zero at the report $q = p$.

is common in machine learning, the scoring rule can be rephrased as a *loss function* by taking the negative, $\ell(r, y) = \|r - y\|_2^2$.

In both the classic scoring rule and the real-valued property elicitation setting, we have general characterizations of which scoring rules are proper and how they can be generated, relating to convex functions [Gneiting and Raftery 2007; Osband 1985; Lambert et al. 2008; Lambert 2018; Steinwart et al. 2014]. It turns out that some properties, such as the variance, cannot be directly elicited by any scoring rule; for example, for a property to be elicitable, its *level sets* $\{p : \Gamma(p) = r\}$, i.e., the set of distributions sharing a particular a property value, must be a convex set. When a property is not elicitable, one often considers “indirect” elicitation, where the agent provides some other information from which one can compute the property of interest. How much other information is required to indirectly elicit a property Γ is known as its *elicitation complexity* [Lambert et al. 2008; Fissler et al. 2016; Frongillo and Kash 2021a]. For example, the variance has elicitation complexity 2, because it can be computed from the mean $\mathbb{E}Y$ and the second moment $\mathbb{E}Y^2$, both of which are elicitable.

1.3 Multi-agent elicitation

There is an extensive literature on eliciting information from groups of agents, often utilizing proper scoring rules as a key tool. Prediction markets [Hanson 2003; Chen and Pennock 2007; Abernethy et al. 2013; Frongillo and Waggoner 2018] and wagering mechanisms [Lambert et al. 2008; Lambert et al. 2015; Freeman and Pennock 2018] are two of the most common. In a prediction market, agents arrive dynamically and make updates to a consensus forecast (equivalently, purchase shares corresponding to predicted events). In wagering mechanisms, agents simultaneously submit “sealed-bid” predictions and wagers. In each case, final payoffs are assigned based on the eventual observed outcome, according to a function typically based on a proper scoring rule. When ground truth is not available—that is, when the outcome y cannot be observed—information can still be elicited by comparing agent reports against each other. This large area of research is referred to as *information elicitation without verification* or the *peer prediction* literature after an eponymous paper [Miller et al. 2005].

2. SINGLE-AGENT ELICITATION

2.1 Elicitability

The term “elicitability”, common in statistics and finance, refers to understanding whether or not a particular property/statistic is elicitable, and if not, what its elicitation complexity is. The fact that elicitable properties have convex level sets has been the main tool to rule out elicibility, such as for the variance and many financial risk measures of interest (see below). Yet since the beginning, one common statistic stood out: the mode, meaning $\Gamma(p) = \operatorname{argmax}_y p(y)$ when p is a probability density function (and suitable generalizations). The mode is interesting because it was widely thought not to be elicitable, yet it does have convex level sets: mixing two distributions with mode r gives a distribution with mode r . Heinrich [2014] showed that indeed the mode is not elicitable, even for some restricted classes of distributions. Even worse, its elicitation complexity is infinite [Dearborn and Frongillo

2020]. More recent work shows that the mode is “asymptotically elicitable”, as the limit of the midpoint of modal intervals [Dimitriadis et al. 2019], and that the mode still fails to be elicitable even when restricting to *strongly unimodal* densities, which have a unique local maximum [Heinrich-Mertsching and Fissler 2022].

A similar story has unfolded in the literature on financial risk measures. Gneiting [2011] caught the attention of this community by observing that Expected Shortfall (ES), a popular risk measure for the regulation of banks, is not elicitable [Embrechts et al. 2014]. A few years later, Fissler and Ziegel [2016] proved that the pair (VaR, ES) is elicitable, where VaR (Value at Risk) is simply a quantile. Their result shows that the elicitation complexity of ES is at most 2; a corresponding lower bound was shown by Frongillo and Kash [2021a]. These results are considered positive, in that they enable “backtesting” procedures to verify that banks are holding enough capital in reserve to compensate for their risk [Fissler et al. 2016]. In parallel to this study of ES, there has been a flurry of research in the statistics and finance communities on the elicibility of various risk and uncertainty measures [Wang and Ziegel 2015; Wang and Wei 2018; Fissler and Ziegel 2021; Fissler et al. 2024; Fissler and Ziegel 2021; Fissler et al. 2021].

Various extensions or generalizations of elicibility have also appeared, such as the “asymptotic elicibility” above. Another notion is *multi-observation elicitation* where one assumes access to multiple independent copies of the outcome Y [Casalaina-Martin et al. 2017; Frongillo et al. 2019]; an example result is that the squared 2-norm of a distribution $\|p\|_2^2$ is elicitable with 2 observations, despite having large elicitation complexity in the usual setting with 1 observation. Finally, the notion of *conditional elicitation* allows one to first elicit a property Γ_1 , and then elicit Γ_2 assuming knowledge of $\Gamma_1(p)$ [Emmer et al. 2015; Frongillo and Kash 2015b; Fissler and Hoga 2024].

A few open problems stand out. First, perhaps the main question remaining in the study of the mode is its elicitation complexity with respect to strongly unimodal distributions. Second, there are many financial risk measures whose elicitation complexity is unknown. One interesting example is the Gini coefficient, given by $(\mathbb{E}|Y_1 - Y_2|)/(2\mathbb{E}Y)$, where Y_1, Y_2 are independent copies of Y [Bellini et al. 2022]. (We note that this property is elicitable with 2 observations, using a scoring rule for the ratio of expectations.) Finally, a nice step toward a general characterization of elicitable vector-valued properties [Frongillo and Kash 2015a] would be to understand the elicitation complexity of the n th central moment, $\Gamma_n(p) = \mathbb{E}_p(Y - \mathbb{E}_p Y)^n$. The main tools for elicitation complexity lower bounds fail for this example, since it is 2-identifiable (a first-order condition) yet the best known upper bound is n , by eliciting the first n moments. The fact that Γ_n is conditionally elicitable conditioned on the mean allows for some partial progress [Frongillo and Kash 2015b].

2.2 Incentives for acquiring information or exerting effort

While proper scoring rules incentivize agents to provide information they already have, a natural question is how to incentivize acquisition of new information, or exertion of effort to produce more accurate predictions. A proper scoring rule will generally pay agents more (in expectation) for better information, so a common approach is to consider the optimal shape of such a scoring rule. A significant amount of recent work has considered aspects of this problem. This work includes Neyman et.

al [2021]; Li et. al [2022] and Hartline et. al [2023] with a motivation of incentivizing effort on the part of e.g. students to learn material; Li and Libgober [2023]; Chen and Yu [2023]; Carrol [2019] and Papireddygar and Waggoner [2022] in the context of contracts; Zhang and Schoenebeck [2023] in a peer-prediction context; and Schoenebeck et. al [2021] in a prediction-markets context.

3. MULTI-AGENT ELICITATION

3.1 Aggregation of forecasts

Given a set of forecasts, how should they be aggregated into a single prediction? This problem has received significant recent theoretical attention, starting with Arieli et. al [2018]. Although the problem does not necessarily involve incentives, it arises naturally in conjunction with e.g. wagering mechanisms and forecasting competitions, which elicit such a set of forecasts.

In the aggregation problem, a set of signals $(S_1, \dots, S_n, Y)$ are drawn jointly from a prior distribution. Here Y is the outcome to be predicted, a binary or real-valued random variable. Each expert $i \in \{1, \dots, n\}$ observes the realization of their signal S_i and updates to a posterior belief with updated expectation $X_i = \mathbb{E}[Y \mid S_i]$. The Bayes-optimal prediction would be $R^* = \mathbb{E}[Y \mid S_1, \dots, S_n]$, the Bayesian aggregation of the information available to all agents. An aggregator collects the individual predictions and produces an estimate $R = R(X_1, \dots, X_n)$. Performance is typically measured as the difference in expected squared loss compared to the optimal aggregator, i.e. $\mathbb{E}[(R - Y)^2 - (R^* - Y)^2]$.

While Arieli et. al [2018] used an additive regret notion to measure performance, Neyman and Roughgarden [2022] considered a competitive-ratio approach. In each case, it is generally impossible to give nontrivial results for arbitrary information structures,² so research focuses on classes of information structures for which positive results are possible. Arieli et. al [2018] considers structures such as conditionally independent signals and Blackwell-ordered experts; Neyman and Roughgarden [2022] considers, in particular, a substitutes condition. Follow-up and other recent work, which generally focuses on improving the bounds for different classes of structures, includes [Levy and Razin 2021; 2022; De Oliveira et al. 2021; Lin and Chen 2023; Guo et al. 2024].

3.2 Information elicitation without verification

The IEWV or “peer prediction” literature is large and generally well-known to the Economics and Computation community; one recent survey in the area is Faltings [2023]. This brief section will be incomplete, but we highlight some important recent progress and interesting recent ideas.

Since Dasgupta and Ghosh [2013], an important paradigm has been the *multi-task* setting in which a group of agents are each asked for their answers to a set of questions, with rewards determined by comparing their sets of answers. Recently, Kong [2020; 2024] gave the *Determinant Mutual Information (DMI)* mechanism which achieves the strongest possible notion of truthfulness using a constant number

²For example, even averaging the agents’ predictions can be an arbitrarily bad idea as compared to e.g. selecting one of them at random.

of questions.³ Zheng et al. [2021] uses theory of property elicitation to give lower bounds (impossibility results) for multi-task elicitation mechanisms.

In the even more challenging *single-question* setting, Schoenebeck and Yu [2023] give strongly-truthful mechanisms with just three agents, each of which answers just a single question. The mechanism is inspired by the Bayesian truth serum of Prelec [2004]; a similar result is given independently by Prelec [2021].

Of other recent work on peer prediction with an elicitation focus, we note Wang et al. [2021], which considers aggregation of forecasts via a peer-prediction style perspective; Liu et al. [2023], which designs “surrogate scoring rules” which can, in some settings, use peer reports to replicate the incentives of a proper scoring rule; and Zhang and Schoenebeck [2023], which considers incentives in peer prediction to exert effort or acquire information.

3.3 Forecasting competitions

In a forecasting competition, a group of agents provide forecasts for a set of events, then the outcome is observed and a single winner is selected. (A *wagering mechanism*, in contrast, may assign various levels of payouts to any or all of the participants.) The Kaggle machine-learning platform, for example, implements forecasting competitions. The “winner-take-all” structure of such competitions distorts the incentives for accurate forecasting. A truthful mechanism was given by Witkowski, et al. [2018], with more results provided in the journal version [Witkowski et al. 2023]. Frongillo, et al. [2021] give a mechanism that is only approximately truthful, but can select the most accurate forecaster with a smaller set of forecasted events. Guarantees have also recently been extended to the setting of correlated events [Frongillo et al. 2023].

3.4 Prediction markets and decentralized finance

The design of automated market makers, algorithms that offer prices to buy or sell assets, originated in the prediction market community as a solution to thin market problems [Hanson 2003]. There has been a significant literature on the theory of prediction markets (see e.g. [Abernethy et al. 2013; Frongillo and Waggoner 2018] and references therein) and efficient implementation in combinatorial settings (e.g. [Dudík et al. 2013; Wang et al. 2021] and references therein). One recent work, Schoenebeck, et al. [2021], considers incentives in prediction markets for exerting effort and acquiring information. Kong and Schoenebeck [2023] examines when information in prediction markets is fully aggregated, relating to when signals are substitutes; Frongillo, et al. [2023] considers a similar question in the case of general communication protocols inspired by prediction markets.

Despite the large literature on the design of automated market makers for prediction markets, their adoption in practice has been limited. By contrast, automated market makers are quite popular mechanisms to run decentralized exchanges in blockchain settings, having traded billions of dollars of assets in recent years [Angeris and Chitra 2020; Angeris et al. 2022]. The dominant paradigm in the blockchain context is the class of *constant-function market makers* (CFMMs). It turns out that CFMMs, while not designed to elicit information per se, are equiv-

³Details of the setting, such as i.i.d. assumptions, are of interest but omitted here.

alent in a strong sense to prediction markets [Frongillo et al. 2024]. We expect this connection to spark a lively exchange of ideas between these two previously independent literatures.

3.5 Mechanism design and fair division

While information elicitation concerns incentives for an agent or group of agents to report honest information, it has deep connections to mechanism design and more generally the problem of allocating goods or services. These connections were perhaps first observed by Fiat, et al. [2013] and later expanded by Frongillo and Kash [2014; 2021b]; these papers give constructions to convert scoring rules to single-agent mechanisms and vice versa, among other connections and results.

The points of contact between these literatures continue to broaden. Particularly emblematic is the recent work showing the equivalence between wagering mechanisms and fair division mechanisms [Freeman et al. 2019; Freeman et al. 2023]. The authors apply this equivalence to a popular wagering mechanism to give the first nontrivial mechanism which is incentive-compatible, proportional, and envy-free. We anticipate more results of this type for multi-agent mechanism design, as researchers continue to leverage the perspective and techniques of information elicitation.

3.6 Online learning from strategic experts

A recent line of work, initiated by [Roughgarden and Schrijvers 2017], asks how to conduct online no-regret learning from expert advice when the experts have incentives to be “chosen” and may misreport their predictions. Exactly truthful mechanisms, built on connections to forecasting competitions, are studied by [Freeman et al. 2020; Mortazavi et al. 2024]. Mechanisms that are only approximately truthful but satisfy good regret guarantees are studied by [Frongillo et al. 2021] and then in the “ m -experts” setting by [Sadeghi and Fazel 2023].

4. CONNECTIONS TO MACHINE LEARNING

The design of scoring rules or loss functions is also an active area of research in machine learning. We first review the basic framework of supervised machine learning, and the role of property elicitation.

4.1 Indirect elicitation in machine learning

In supervised machine learning, we wish to learn a model or hypothesis which makes a prediction given some feature vector $x \in \mathcal{X}$. Most algorithms employ *empirical risk minimization (ERM)*, which simply chooses a hypothesis $h : \mathcal{X} \rightarrow \mathbb{R}^d$ from a class $\mathcal{H}$ that minimizes loss over a data set:

$$h^* \in \operatorname{argmin}_{h \in \mathcal{H}} \sum_{(x,y) \in \text{data}} L(h(x), y) . \quad (1)$$

For example, in ordinary least-squares (OLS) regression, $\mathcal{H}$ consists of linear functions from $\mathcal{X}$ to $\mathbb{R}$, and $L_2(r, y) = (r - y)^2$ is squared error. In regression problems, the statistical consistency (“correctness”) of ERM boils down to a question of property elicitation for the conditional distributions $\Pr[Y|X=x]$: whether L elicits the desired conditional statistic (Fig. 1(a)). As squared error elicits the mean, OLS

therefore fits to the conditional mean $h^*(x) = \mathbb{E}[Y|X=x]$ so long as $h^* \in \mathcal{H}$ (see below). Similarly, absolute loss $L_1(r, y) = |r - y|$ yields median regression.

In discrete prediction problems such as classification, ranking, and structured prediction, we are instead given a target discrete loss $\ell : \mathcal{R} \times \mathcal{Y} \rightarrow \mathbb{R}$ for finite sets $\mathcal{R}$ (predictions) and $\mathcal{Y}$ (labels), and wish to learn a hypothesis $h_{\text{targ}} : \mathcal{X} \rightarrow \mathcal{R}$ achieving low expected loss $\mathbb{E}_D \ell(h_{\text{targ}}(X), Y)$ over the underlying distribution D . For example, traditional classification has $\mathcal{R} = \mathcal{Y}$ with $\ell(r, y) = \mathbb{1}\{r \neq y\}$ being 0-1 loss (penalty 0 if correct, 1 if incorrect). Solving ERM (1) is generally NP-hard for discrete losses ℓ , so instead we seek a *surrogate* loss $L : \mathbb{R}^d \times \mathcal{Y} \rightarrow \mathbb{R}$ which is convex in the first argument. This d is called the *prediction dimension*, and plays a key role in structured prediction, where it can be exponentially large in the natural dimension of the problem [Osokin et al. 2017]. The hypothesis $h_{\text{surg}} : \mathcal{X} \rightarrow \mathbb{R}^d$ is then converted to one answering the target problem via a *link function* $\psi : \mathbb{R}^d \rightarrow \mathcal{R}$ mapping back to target predictions.

$$h_{\text{surg}} \in \operatorname{argmin}_{h \in \mathcal{H}} \sum_{(x,y) \in \text{data}} L(h(x), y), \quad h_{\text{targ}} = \psi \circ h_{\text{surg}}. \quad (2)$$

Many algorithms follow this paradigm, including support vector machines (SVMs), logistic regression, boosting, and deep neural networks.

For surrogate ERM (2), consistency means that low surrogate (L) loss of h_{surg} should imply low target (ℓ) loss of the linked hypothesis $h_{\text{targ}} = \psi \circ h_{\text{surg}}$ given more and more data. Consistency is a precursor to *rates*, which quantify how fast target loss is minimized. When the class $\mathcal{H}$ is sufficiently rich, consistency of surrogate minimization reduces to *calibration*, a condition stating that predictions cannot approach the optimal surrogate loss while linking to incorrect target predictions [Bartlett et al. 2006; Tewari and Bartlett 2007; Agarwal and Agarwal 2015]. Just as with regression problems, calibration only depends on the loss with respect to conditional distributions on $\mathcal{Y}$. Moreover, calibration implies “in-




Fig. 1. For rich hypothesis classes $\mathcal{H}$, consistency of a learning algorithm reduces to its behavior on the conditional distributions $Y|X=x$ for each $x \in \mathcal{X}$. **(a)** In linear regression, minimizing absolute loss L_1 yields hypothesis h_1^* , where $h_1^*(x)$ tracks the median of $Y|X=x$. Squared loss L_2 instead tracks the conditional mean. **(b)** In SVMs for classification, minimizing hinge loss conditioned on $X=x$, and taking the sign of the result (the link function), will minimize conditional 0-1 loss.

direct” property elicitation, meaning $\gamma = \psi \circ \Gamma$ where ℓ elicits γ and L elicits Γ .⁴ For example, 0-1 loss elicits the mode (most likely label), and the SVM hinge loss $L_{\text{hinge}}(r, y) = \max(0, 1 - ry)$ indirectly elicits the mode with the link $\psi(r) = \text{sign}(r)$ (Fig. 1(b)), where here $r \in \mathbb{R}$ and $y \in \mathcal{Y} = \{-1, 1\}$.

4.2 Surrogate loss design

Compared to the property elicitation results discussed in § 2.1, perhaps the most significant constraint in machine learning settings is that the surrogate loss should ideally be convex and thus efficient to optimize. Somewhat surprisingly, convexity of the loss comes for free for continuous real-valued properties [Finocchiario and Frongillo 2018], i.e., in prediction dimension 1, but for higher prediction dimensions this is no longer the case. By analogy to elicitation complexity, recent work has tried to understand the lowest prediction dimension needed for a consistent surrogate loss to exist [Agarwal and Agarwal 2015; Ramaswamy and Agarwal 2016; Finocchiario et al. 2020]. One interesting example is the *abstain property*, which takes the value of abstain ($\perp$) if no label has probability at least 0.5, and is the most likely label otherwise. Naively a convex loss for this property would require prediction dimension $n - 1$ for n labels, but a clever construction due to Ramaswamy, et al. [2018] uses only $\log n$ dimensions. Motivated by this work, Finocchiario et al. [2019; 2024] give a general framework to design consistent convex surrogate loss functions for any target. Despite this progress, however, many important questions remain; perhaps most glaring is the lack of tools to bound the prediction dimension required for most target problems.

4.3 Multicalibration and decision robustness

A particular trend of interest is the problem of predicting without knowing which scoring rule (or loss function) one is predicting for. Constructing such “omnipredictors” is the focus of Gopalan et. al [2022] and Hu et. al [2023]. We are given a dataset and, instead of a single target loss, a family of loss functions $\mathcal{L}$. The goal is to learn a hypothesis h that performs well on every loss in $\mathcal{L}$ simultaneously, possibly with loss-specific post-processing. A conceptually similar problem is studied in an online learning setting by Kleinberg et. al [2023]; see also Ehm et al. [2016].

These works are closely related to *multicalibration* [Hebert-Johnson et al. 2018], a concept of interest in fair machine learning. In the binary classification setting, a hypothesis $h : \mathcal{X} \rightarrow [0, 1]$ is called calibrated if, among the subset of pairs (x, y) for which $h(x) = c$, a c -fraction have true label $y = 1$. In multicalibration, h must be calibrated even when conditioning on particular subgroups; this turns out to be useful for achieving omniprediction. In the spirit of property elicitation, Jung et. al [2021] extends multicalibration from means to properties such as the variance.

Acknowledgements and requests for suggestions

We thank Tilmann Gneiting, Ian Kash, and Jens Witkowski. We appreciate and welcome any additional comments or suggestions of works that we have missed.

⁴It remains an interesting open problem to study exactly when and how calibration and indirect elicitation differ in practical examples.

REFERENCES

- ABERNETHY, J., CHEN, Y., AND VAUGHAN, J. W. 2013. Efficient market making via convex optimization, and a connection to online learning. *ACM Transactions on Economics and Computation* 1, 2, 12.
- AGARWAL, A. AND AGARWAL, S. 2015. On consistent surrogate risk minimization and property elicitation. In *Conference on Learning Theory*. 4–22.
- ANGERIS, G., AGRAWAL, A., EVANS, A., CHITRA, T., AND BOYD, S. 2022. Constant function market makers: Multi-asset trades via convex optimization. In *Handbook on Blockchain*. Springer, 415–444.
- ANGERIS, G. AND CHITRA, T. 2020. Improved price oracles: Constant function market makers. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*. 80–91.
- ARIELI, I., BABICHENKO, Y., AND SMORODINSKY, R. 2018. Robust forecast aggregation. *Proceedings of the National Academy of Sciences* 115, 52, E12135–E12143.
- BARTLETT, P. L., JORDAN, M. I., AND MCAULIFFE, J. D. 2006. Convexity, classification, and risk bounds. *Journal of the American Statistical Association* 101, 473, 138–156.
- BELLINI, F., FADINA, T., WANG, R., AND WEI, Y. 2022. Parametric measures of variability induced by risk measures. *Insurance: Mathematics and Economics* 106, 270–284.
- BRIER, G. 1950. Verification of forecasts expressed in terms of probability. *Monthly weather review* 78, 1, 1–3.
- CARROLL, G. 2019. Robust incentives for information acquisition. *Journal of Economic Theory* 181, 382–420.
- CASALAINA-MARTIN, S., FRONGILLO, R., MORGAN, T., AND WAGGONER, B. 2017. Multi-observation elicitation. In *Conference on Learning Theory*. PMLR, 449–464.
- CHEN, Y. AND PENNOCK, D. 2007. A utility framework for bounded-loss market makers. In *Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence*. 49–56.
- CHEN, Y. AND YU, F.-Y. 2023. Optimal scoring rule design under partial knowledge.
- DASGUPTA, A. AND GHOSH, A. 2013. Crowdsourced judgement elicitation with endogenous proficiency. In *Proceedings of the 22nd International Conference on World Wide Web*. 319–330.
- DE OLIVEIRA, H., ISHII, Y., AND LIN, X. 2021. Robust merging of information. In *Proceedings of the 22nd ACM Conference on Economics and Computation*. EC '21. Association for Computing Machinery, New York, NY, USA, 341–342.
- DEARBORN, K. AND FRONGILLO, R. 2020. On the indirect elicibility of the mode and modal interval. *Annals of the Institute of Statistical Mathematics* 72, 5, 1095–1108.
- DIMITRIADIS, T., PATTON, A. J., AND SCHMIDT, P. W. 2019. Testing forecast rationality for measures of central tendency. *arXiv preprint arXiv:1910.12545*.
- DUDÍK, M., LAHAIE, S., PENNOCK, D. M., AND ROTHSCHILD, D. 2013. A combinatorial prediction market for the us elections. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*. 341–358.
- EHM, W., GNEITING, T., JORDAN, A., AND KRÜGER, F. 2016. Of quantiles and expectiles: consistent scoring functions, choquet representations and forecast rankings. *Journal of the Royal Statistical Society Series B: Statistical Methodology* 78, 3, 505–562.
- EMBRECHTS, P., PUCETTI, G., RÜSCHENDORF, L., WANG, R., AND BELERAJ, A. 2014. An Academic Response to Basel 3.5. *Risks* 2, 1 (Feb.), 25–48.
- EMMER, S., KRATZ, M., AND TASCHE, D. 2015. What is the best risk measure in practice? a comparison of standard measures. *Journal of Risk* 18, 2, 31–60.
- FALTINGS, B. 2023. Game-theoretic mechanisms for eliciting accurate information. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*. IJCAI '23.
- FIAT, A., KARLIN, A., KOUTSOUPAS, E., AND VIDALI, A. 2013. Approaching Utopia: strong truthfulness and externality-resistant mechanisms. In *Proceedings of the 4th conference on Innovations in Theoretical Computer Science*. 221–230.
- FINOCCHIARO, J. AND FRONGILLO, R. 2018. Convex elicitation of continuous properties. In *Advances in Neural Information Processing Systems* 31, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, Eds. Curran Associates, Inc., 10425–10434.

- FINOCCHIARO, J., FRONGILLO, R., AND WAGGONER, B. 2019. An embedding framework for consistent polyhedral surrogates. In *Proceedings of Advances In Neural Information Processing Systems (NeurIPS)*. 10780–10790.
- FINOCCHIARO, J., FRONGILLO, R., AND WAGGONER, B. 2020. Embedding dimension of polyhedral losses. In *Proceedings of Conference on Learning Theory (COLT)*.
- FINOCCHIARO, J., FRONGILLO, R. M., AND WAGGONER, B. 2024. An embedding framework for the design and analysis of consistent polyhedral surrogates. *Journal of Machine Learning Research* 25, 63, 1–60.
- FISSLER, T., FRONGILLO, R., HLAVINOVÁ, J., AND RUDLOFF, B. 2021. Forecast evaluation of quantiles, prediction intervals, and other set-valued functionals. *Electronic Journal of Statistics* 15, 1034–1084.
- FISSLER, T. AND HOGA, Y. 2024. Backtesting systemic risk forecasts using multi-objective elicibility. *Journal of Business & Economic Statistics* 42, 2, 485–498.
- FISSLER, T., LIU, F., WANG, R., AND WEI, L. 2024. Elicitability and identifiability of tail risk measures. *arXiv preprint arXiv:2404.14136*.
- FISSLER, T., ZIEGEL, J., AND GNEITING, T. 2016. Expected Shortfall is jointly elicitable with Value at Risk-Implications for backtesting. *Risk Magazine*.
- FISSLER, T. AND ZIEGEL, J. F. 2016. Higher order elicibility and Osband’s principle. *The Annals of Statistics* 44, 4, 1680–1707.
- FISSLER, T. AND ZIEGEL, J. F. 2021. On the elicibility of range value at risk. *Statistics & risk modeling* 38, 1-2, 25–46.
- FREEMAN, R., PENNOCK, D., PODIMATA, C., AND VAUGHAN, J. W. 2020. No-regret and incentive-compatible online learning. In *International Conference on Machine Learning*. PMLR, 3270–3279.
- FREEMAN, R. AND PENNOCK, D. M. 2018. An axiomatic view of the parimutuel consensus wagering mechanism. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 1936–1938.
- FREEMAN, R., PENNOCK, D. M., AND VAUGHAN, J. W. 2019. An equivalence between wagering and fair-division mechanisms. In *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 33. 1957–1964.
- FREEMAN, R., WITKOWSKI, J., WORTMAN VAUGHAN, J., AND PENNOCK, D. M. 2023. An equivalence between fair division and wagering mechanisms. *Management Science* 0, 0.
- FRONGILLO, R., GOMEZ, R., THILAGAR, A., AND WAGGONER, B. 2021. Efficient competitions and online learning with strategic forecasters. In *Proceedings of the 22nd ACM Conference on Economics and Computation*. 479–496.
- FRONGILLO, R. AND KASH, I. 2014. General truthfulness characterizations via convex analysis. In *Web and Internet Economics*. Springer, 354–370.
- FRONGILLO, R. AND KASH, I. 2015a. Vector-Valued Property Elicitation. In *Proceedings of the 28th Conference on Learning Theory*. 1–18.
- FRONGILLO, R. AND KASH, I. A. 2015b. On Elicitation Complexity and Conditional Elicitation. *arXiv preprint arXiv:1506.07212*.
- FRONGILLO, R., LLADSER, M., THILAGAR, A., AND WAGGONER, B. 2023. Forecasting competitions with correlated events. *arXiv preprint arXiv:2303.13793*.
- FRONGILLO, R., MEHTA, N. A., MORGAN, T., AND WAGGONER, B. 2019. Multi-Observation Regression. In *The 22nd International Conference on Artificial Intelligence and Statistics*. 2691–2700.
- FRONGILLO, R., NEYMAN, E., AND WAGGONER, B. 2023. Agreement implies accuracy for substitutable signals. In *Proceedings of the 2023 ACM Conference on Economics and Computation*. EC. ACM.
- FRONGILLO, R., PAPIREDDYGARI, M., AND WAGGONER, B. 2024. An Axiomatic Characterization of CFMMs and Equivalence to Prediction Markets. In *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*. Vol. 287. 51:1–51:21.
- FRONGILLO, R. AND WAGGONER, B. 2018. An axiomatic study of scoring rule markets. In *Proceedings of the 9th Conference on Innovations in Theoretical Computer Science (ITCS)*. 15:1–15:20.

- FRONGILLO, R. M. AND KASH, I. A. 2021a. Elicitation complexity of statistical properties. *Biometrika* 108, 4, 857–879.
- FRONGILLO, R. M. AND KASH, I. A. 2021b. General truthfulness characterizations via convex analysis. *Games and Economic Behavior* 130, 636–662.
- GNEITING, T. 2011. Making and Evaluating Point Forecasts. *Journal of the American Statistical Association* 106, 494, 746–762.
- GNEITING, T. AND RAFTERY, A. E. 2007. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association* 102, 477, 359–378.
- GOOD, I. J. 1952. Rational decisions. *Journal of the Royal Statistical Society. Series B (Methodological)*, 107–114.
- GOPALAN, P., KALAI, A. T., REINGOLD, O., SHARAN, V., AND WIEDER, U. 2022. Omnipredictors. In *13th Innovations in Theoretical Computer Science Conference (ITCS 2022)*, M. Braverman, Ed. Leibniz International Proceedings in Informatics (LIPIcs), vol. 215. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 79:1–79:21.
- GUO, Y., HARTLINE, J. D., HUANG, Z., KONG, Y., SHAH, A., AND YU, F.-Y. 2024. Algorithmic robust forecast aggregation.
- HANSON, R. 2003. Combinatorial Information Market Design. *Information Systems Frontiers* 5, 1, 107–119.
- HARTLINE, J. D., SHAN, L., LI, Y., AND WU, Y. 2023. Optimal scoring rules for multi-dimensional effort. In *Proceedings of Thirty Sixth Conference on Learning Theory*, G. Neu and L. Rosasco, Eds. Proceedings of Machine Learning Research, vol. 195. PMLR, 2624–2650.
- HEBERT-JOHNSON, U., KIM, M., REINGOLD, O., AND ROTHBLUM, G. 2018. Multicalibration: Calibration for the (Computationally-identifiable) masses. In *Proceedings of the 35th International Conference on Machine Learning*, J. Dy and A. Krause, Eds. Proceedings of Machine Learning Research, vol. 80. PMLR, 1939–1948.
- HEINRICH, C. 2014. The mode functional is not elicitable. *Biometrika* 101, 1, 245–251.
- HEINRICH-MERTSCHING, C. AND FISSLER, T. 2022. Is the mode elicitable relative to unimodal distributions? *Biometrika* 109, 4, 1157–1164.
- HU, L., NAVON, I. R. L., REINGOLD, O., AND YANG, C. 2023. Omnipredictors for constrained optimization. In *International Conference on Machine Learning*. PMLR, 13497–13527.
- JUNG, C., LEE, C., PAI, M., ROTH, A., AND VOHRA, R. 2021. Moment multicalibration for uncertainty estimation. In *Conference on Learning Theory*. PMLR, 2634–2678.
- KLEINBERG, B., LEME, R. P., SCHNEIDER, J., AND TENG, Y. 2023. U-calibration: Forecasting for an unknown agent. In *The Thirty Sixth Annual Conference on Learning Theory*. PMLR, 5143–5145.
- KONG, Y. 2020. Dominantly truthful multi-task peer prediction with a constant number of tasks. In *Proceedings of the Fourteenth annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2398–2411.
- KONG, Y. 2024. Dominantly truthful peer prediction mechanisms with a finite number of tasks. *J. ACM* 71, 2 (apr).
- KONG, Y. AND SCHOENEBECK, G. 2023. False Consensus, Information Theory, and Prediction Markets. In *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*, Y. Tauman Kalai, Ed. Leibniz International Proceedings in Informatics (LIPIcs), vol. 251. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 81:1–81:23.
- LAMBERT, N. 2018. Elicitation and Evaluation of Statistical Forecasts. *Preprint*.
- LAMBERT, N. S., LANGFORD, J., VAUGHAN, J. W., CHEN, Y., REEVES, D. M., SHOHAM, Y., AND PENNOCK, D. M. 2015. An axiomatic characterization of wagering mechanisms. *Journal of Economic Theory* 156, 389–416.
- LAMBERT, N. S., LANGFORD, J., WORTMAN, J., CHEN, Y., REEVES, D., SHOHAM, Y., AND PENNOCK, D. 2008. Self-financed wagering mechanisms for forecasting. In *Proceedings of the 9th ACM conference on Electronic commerce*. ACM, 170–179.
- LAMBERT, N. S., PENNOCK, D. M., AND SHOHAM, Y. 2008. Eliciting properties of probability distributions. In *Proceedings of the 9th ACM Conference on Electronic Commerce*. 129–138.

- LEVY, G. AND RAZIN, R. 2021. A maximum likelihood approach to combining forecasts. *Theoretical Economics* 16, 1, 49–71.
- LEVY, G. AND RAZIN, R. 2022. Combining forecasts in the presence of ambiguity over correlation structures. *Journal of Economic Theory* 199, 105075. Symposium Issue on Ambiguity, Robustness, and Model Uncertainty.
- LI, Y., HARTLINE, J. D., SHAN, L., AND WU, Y. 2022. Optimization of scoring rules. In *Proceedings of the 23rd ACM Conference on Economics and Computation*. EC '22. Association for Computing Machinery, New York, NY, USA, 988–989.
- LI, Y. AND LIBGOBER, J. 2023. Optimal scoring for dynamic information acquisition.
- LIN, T. AND CHEN, Y. 2023. Sample complexity of forecast aggregation. In *Advances in Neural Information Processing Systems*, A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, Eds. Vol. 36. Curran Associates, Inc., 36053–36093.
- LIU, Y., WANG, J., AND CHEN, Y. 2023. Surrogate scoring rules. *ACM Trans. Econ. Comput.* 10, 3 (feb).
- MILLER, N., RESNICK, P., AND ZECKHAUSER, R. 2005. Eliciting informative feedback: The peer-prediction method. *Management Science* 51, 9, 1359–1373.
- MORTAZAVI, A., LIN, J., AND MEHTA, N. 2024. On the price of exact truthfulness in incentive-compatible online learning with bandit feedback: a regret lower bound for WSU-UX. In *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, S. Dasgupta, S. Mandt, and Y. Li, Eds. Proceedings of Machine Learning Research, vol. 238. PMLR, 4681–4689.
- NEYMAN, E., NOAROV, G., AND WEINBERG, S. M. 2021. Binary scoring rules that incentivize precision. In *Proceedings of the 22nd ACM Conference on Economics and Computation*. 718–733.
- NEYMAN, E. AND ROUGHGARDEN, T. 2022. Are you smarter than a random expert? the robust aggregation of substitutable signals. In *Proceedings of the 23rd ACM Conference on Economics and Computation*. EC '22. Association for Computing Machinery, New York, NY, USA, 990–1012.
- OSBAND, K. H. 1985. *Providing Incentives for Better Cost Forecasting*. University of California, Berkeley.
- OSOKIN, A., BACH, F., AND LACOSTE-JULIEN, S. 2017. On structured prediction theory with calibrated convex surrogate losses. In *Advances in Neural Information Processing Systems*. 302–313.
- PAPIREDDYGARI, M. AND WAGGONER, B. 2022. Contracts with information acquisition, via scoring rules. In *Proceedings of the 2022 ACM Conference on Economics and Computation*. EC. ACM.
- PRELEC, D. 2004. A Bayesian Truth Serum for Subjective Data. *Science* 306, 5695 (Oct.), 462–466.
- PRELEC, D. 2021. Bilateral bayesian truth serum: The nxm signals case. *Available at SSRN 3908446*.
- RAMASWAMY, H. G. AND AGARWAL, S. 2016. Convex calibration dimension for multiclass loss matrices. *The Journal of Machine Learning Research* 17, 1, 397–441.
- RAMASWAMY, H. G., TEWARI, A., AND AGARWAL, S. 2018. Consistent algorithms for multiclass classification with an abstain option. *Electronic Journal of Statistics* 12, 1, 530–554.
- ROUGHGARDEN, T. AND SCHRIJVERS, O. 2017. Online prediction with selfish experts. *Advances in Neural Information Processing Systems* 30.
- SADEGHI, O. AND FAZEL, M. 2023. No-regret online prediction with strategic experts. In *Advances in Neural Information Processing Systems*, A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, Eds. Vol. 36. Curran Associates, Inc., 54696–54715.
- SAVAGE, L. 1971. Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association* 66, 336, 783–801.
- SCHOENEBECK, G., YU, C., AND YU, F.-Y. 2021. Timely information from prediction markets. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*. AAMAS '21. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1145–1153.

- SCHOENEBECK, G. AND YU, F.-Y. 2023. Two strongly truthful mechanisms for three heterogeneous agents answering one question. *ACM Trans. Econ. Comput.* 10, 4 (feb).
- STEINWART, I., PASIN, C., WILLIAMSON, R., AND ZHANG, S. 2014. Elicitation and Identification of Properties. In *Proceedings of The 27th Conference on Learning Theory*. 482–526.
- TEWARI, A. AND BARTLETT, P. L. 2007. On the consistency of multiclass classification methods. *The Journal of Machine Learning Research* 8, 1007–1025.
- WANG, J., LIU, Y., AND CHEN, Y. 2021. Forecast aggregation via peer prediction. *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing* 9, 1 (Oct.), 131–142.
- WANG, R. AND WEI, Y. 2018. Risk functionals with convex level sets. *Available at SSRN 3292661*.
- WANG, R. AND ZIEGEL, J. F. 2015. Elicitable distortion risk measures: A concise proof. *Statistics & Probability Letters* 100, 172–175.
- WANG, X., PENNOCK, D. M., DEVANUR, N. R., ROTHSCILD, D. M., TAO, B., AND WELLMAN, M. P. 2021. Designing a combinatorial financial options market. In *Proceedings of the 22nd ACM Conference on Economics and Computation*. 864–883.
- WITKOWSKI, J., FREEMAN, R., VAUGHAN, J. W., PENNOCK, D. M., AND KRAUSE, A. 2018. Incentive-compatible forecasting competitions. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*. AAAI.
- WITKOWSKI, J., FREEMAN, R., VAUGHAN, J. W., PENNOCK, D. M., AND KRAUSE, A. 2023. Incentive-compatible forecasting competitions. *Management Science* 69, 3, 1354–1374.
- ZHANG, Y. AND SCHOENEBECK, G. 2023. High-effort crowds: Limited liability via tournaments. In *Proceedings of the ACM Web Conference 2023*. WWW '23. Association for Computing Machinery, New York, NY, USA, 3467–3477.
- ZHENG, S., YU, F.-Y., AND CHEN, Y. 2021. The limits of multi-task peer prediction. In *Proceedings of the 22nd ACM Conference on Economics and Computation*. EC '21. Association for Computing Machinery, New York, NY, USA, 907–926.

Online Matching: A Brief Survey

ZHIYI HUANG

University of Hong Kong

and

ZHILAO GAVIN TANG

Shanghai University of Finance and Economics

and

DAVID WAJC

Technion — Israel Institute of Technology

Matching, capturing allocation of items to unit-demand buyers, or tasks to workers, or pairs of collaborators, is a central problem in economics. Indeed, the growing prevalence of matching-based markets, many of which online in nature, has motivated much research in economics, operations research, computer science, and their intersection. This brief survey is meant as an introduction to the area of online matching, with an emphasis on recent trends, both technical and conceptual.

Categories and Subject Descriptors: F.2.2 [Nonnumerical Algorithms and Problems]: Matching Algorithms

General Terms: Algorithms

Additional Key Words and Phrases: Matching, Online Algorithms, Online Matching, Survey

1. INTRODUCTION

Matching theory lies at the heart of Economics, Computation and their intersection. Matching markets have played increasingly dominant roles in the world economy, both on the micro and macro scale. Such markets arise in domains as varied as Internet advertising, crowdsourcing of work and transportation, and organ transplantation. The repeated interactions and lack of certainty about future participants (buyers, sellers, etc.) result in these industries' dynamics being prime examples of *online* matching markets. See [Echenique et al. 2023] for detailed discussions. In this brief survey, we focus on three main aspects of recent developments in the study of such online matching and allocation problems.

Particularly prevalent are online *bipartite* matching markets. Often, agents on one side of the market arrive up front, while agents on the other side are revealed sequentially, to be matched (or not) immediately and irrevocably. For example, this dynamic abstracts the Internet advertising marketplace, with advertisers known up front and user queries (ad slots) revealed online. This motivates the study of online bipartite matching [Karp et al. 1990], and its generalization to weighted settings [Aggarwal et al. 2011; Feldman et al. 2009] and the AdWords problems [Mehta et al. 2007]. For more on the motivation from the Internet advertising application, see the influential survey on online bipartite matching and ad allocation by [Mehta 2013], and the more recent survey by [Devanur and Mehta 2022]. We outline recent

Authors' addresses: zhiyi@cs.hku.hk, tang.zhihao@mail.shufe.edu.cn, david.wajc@gmail.com

developments on these online bipartite matching and online ad allocation problems, in §2.

There are, of course, many aspects of modern online matching markets that are not bipartite, or that allow for agents on either side of bipartite matching markets to enter in an interleaved order. Similarly, while classic online matching models consider agents as having left the market after matching, in many crowdsourcing marketplaces (e.g., DoorDash, TaskRabbit, Uber/Lyft, etc.) freelance workers return to the market after being assigned a task (i.e., being matched) and completing their tasks. We discuss recent progress on modeling and addressing such problems, in §3.

The above-mentioned sections focus on the robust, but somewhat pessimistic, modeling choice of adversarial inputs and arrival orders. A less pessimistic model is that of random-order arrivals (“secretary models”), where the input is generated adversarially, but permuted by nature. Another modeling choice, motivated by the abundance of historical data from which to learn trends, is to posit a stochastic arrival model with parameters known to the algorithm. Here one can compare with either the best offline algorithm (computed by a “prophet” who knows the future) or the best online algorithm (computed by a “philosopher” who has enough time to think/compute). We discuss progress on online matching for such stochastic models, and their connection to mechanism design, in §4.

Finally, we give a brief glimpse of some overarching techniques that have played key recurring roles in the aforementioned recent developments, in §5. We illustrate some of the ideas with particularly short (and in our opinion, quite teachable) examples of some of the basic techniques in this space, in §A.

2. ONLINE MATCHING AND AD ALLOCATION

Researchers have made much progress in the past decade on the aforementioned online bipartite matching and ad allocation problems.

The most general problem along this line is *online submodular welfare maximization*. Consider a set of offline agents (advertisers), and a set of online items (impressions) that arrive one at a time. Each agent a has a submodular value function v_a over subsets of items; the algorithm can evaluate $v_a(S)$ for any subset S of the arrived items. On arrival of an online item, the algorithm must allocate it to an agent immediately and irrevocably. The basic benchmark is the greedy algorithm that allocates each impression to maximize the immediate increase in social welfare, which is $1/2$ -competitive. For the general problem, this ratio is optimal for polynomial-time algorithms under standard complexity-theoretic assumptions [Kapralov et al. 2013]. Most research has therefore focused on special cases of interest of this problem.

The (*unweighted*) *online bipartite matching* problem is the special case when each agent either likes or dislikes an item, and is willing to pay \$1 to get any one item they like: the value $v_a(S)$ equals 1 if agent a likes at least one item in S , and is 0 otherwise. Further, when a subset of items S is allocated to agent a , we can interpret this as matching a to any one item that a likes in S (e.g., the first one).

In online advertising, some advertisers may be able to pay more than others for an impression that they like. This can be captured by the *vertex-weighted*

generalization of online bipartite matching, where the value $v_a(S)$ is agent a 's weight w_a if agent a likes at least one item in S , and is 0 otherwise.

More generally, the same advertiser may have different values for different impressions, e.g., depending on the users' cookies and other information. This motivates considering edge weights instead of vertex weights; the advertiser only pays for one impression like in the unweighted and vertex-weighted case.¹ This is the *display ads* problem introduced by [Feldman et al. 2009], where $v_a(S)$ equals the largest edge weight w_{ai} among items $i \in S$.

Last but not least, some platforms let advertisers set a daily budget rather than a limit on the number of impressions. Given the allocation of impressions in a day, each advertiser a pays either the sum of its values for impressions it gets or its daily budget B_a , whichever is smaller; in other words, $v_a(S) = \min \{ \sum_{i \in S} w_{ai}, B_a \}$. Here, the weight w_{ai} is often referred to as agent a 's *bid* for impression i . This is the *AdWords* problem introduced by [Mehta et al. 2007].

Given the uncertainty over future items, when we decide how to allocate an online item, we do not want to put all our eggs in the same basket. It is easier to implement this old wisdom when the item is divisible (alternatively, if each agent has a large basket that can take many items). Imagine that each online item carries one litre of water (a divisible egg); each offline agent has a bucket (basket) of capacity one litre. The algorithm distributes an online item's water to its neighbors, where the amount of water going to each neighbor represents the fraction of the item allocated to the agent. The BALANCE algorithm (a.k.a., WATER LEVEL or WATER-FILLING [Azar and Litichevsky 2006]) lets the water flow to the least loaded bucket (the basket with the least amount of eggs); if there are multiple least loaded buckets, the water flows to them at an equal rate. This algorithm and its generalizations achieve the optimal $1 - 1/e$ competitive ratio in all the mentioned special cases of online submodular welfare maximization, including unweighted matching [Kalyanasundaram and Pruhs 2000], vertex-weighted matching [Buchbinder et al. 2007], display ads [Feldman et al. 2009], and AdWords [Mehta et al. 2007].

In the original problems where items are indivisible, we can distribute the risk through randomized decisions. However, making independent random decisions in each round is insufficient for getting a competitive ratio better than $1/2$, the baseline set by the greedy algorithm [Karp et al. 1990]. In the same paper that [Karp et al. 1990] introduced the online bipartite matching problem, they also gave an elegant RANKING algorithm achieving the optimal $1 - 1/e$ competitive ratio, later generalized by [Aggarwal et al. 2011] to vertex-weighted problem. The algorithm can be viewed as letting each offline vertex independently set a random price, and having each online vertex choose the lowest price offered by its unmatched neighbors. See §5 for a further discussion on this economic interpretation of RANKING and other online matching algorithms.

2.1 Breaking the $1/2$ Barrier in Longstanding Open Problems

Recall that for online (vertex-weighted) bipartite matching RANKING achieves an optimal competitive ratio, and in particular breaks the barrier of $1/2$. For display ads and AdWords, however, finding an online algorithm strictly better than the

¹Higher capacity can be simulated by creating multiple offline vertices per advertiser.

$1/2$ -competitive greedy algorithm had remained elusive for more than a decade. Fundamentally new ideas seemed necessary because a critical invariant in the analysis of RANKING fails to hold in these two problems.

In 2020, the $1/2$ barrier was broken for both problems using a new technique called Online Correlated Selection (OCS); see §5.4 for a further discussion on this technique. [Fahrbach et al. 2022] introduced the concept of OCS and gave a 0.508-competitive algorithm for display ads. [Huang et al. 2020] modified the definition of OCS and applied it to the AdWords problem, and as a result, obtained a 0.501-competitive algorithm. The OCS technique has then been improved in a series of works by [Shin and An 2021], [Gao et al. 2021], and [Blanc and Charikar 2021]. The state-of-the-art competitive ratio for display ads is 0.536, given by a multi-way OCS algorithm by [Blanc and Charikar 2021].

Despite the aforementioned progress, we remark that there is no known evidence that the $(1 - 1/e)$ -competitive ratio cannot be achieved in display ads and AdWords. Hence, closing the gaps between the upper and lower bounds for these two problems remains an important open problem.

For the general online submodular welfare maximization problem, we recall that the simple greedy algorithm is $1/2$ competitive, and (barring surprising developments in complexity theory) no polynomial-time online algorithms can do better [Kapralov et al. 2013]. That being said, this impossibility relied on the computational hardness of maximizing a submodular function. It would be interesting to explore online algorithms with unlimited computational capacity, because practical heuristics can often solve these optimization problems better than the worst-case approximation ratio promises, and positive results along this line may point to other special submodular functions that are computationally tractable.

	Fractional/Divisible Relaxation	Original Problem
Unweighted	$1 - 1/e$ [Kalyanasundaram and Pruhs 2000]	$1 - 1/e$ [Karp et al. 1990]
Vertex-Weighted	$1 - 1/e$ [Buchbinder et al. 2007, Section 5]	$1 - 1/e$ [Aggarwal et al. 2011]
Display Ads	$1 - 1/e$ [Feldman et al. 2009]	0.536 [Blanc and Charikar 2021]
AdWords	$1 - 1/e$ [Mehta et al. 2007]	0.501 [Huang et al. 2020]

Table I. State-of-the-art for central online bipartite matching & allocation problems.

2.2 Stochastic Rewards and Oblivious Budget

Many online advertising platforms adopt the *pay-per-click* model. In this model, an advertiser pays each time a user clicks on its advertisement. Since these platforms cannot control the user's behavior, they resort to the next best option: modeling a user's behavior stochastically, and estimating the probability that the user clicks the advertisement, known as the *click-through-rate* (CTR). As a result of the users' stochastic behavior, the platform's revenue from assigning an impression to an advertiser is also stochastic.

[Mehta and Panigrahi 2012] introduced online matching with stochastic rewards. They analyzed both the RANKING algorithm and a variant of the BALANCE algorithm, and showed competitive ratios better than $1/2$ for uniform CTRs. That is, the CTR is either p or 0 (if the advertiser is not interested in this impression).

Progress made in the past decade on stochastic rewards is threefold. First, [Mehta et al. 2014] gave the first algorithm that breaks the $1/2$ barrier for non-uniform but sufficiently small CTRs. On one hand, small CTRs are arguably the most relevant case in practice because most keywords' CTRs are less than 10%. On the other hand, 10% or even 1% is larger than the assumption made by [Mehta et al. 2014] and its follow-up works. Hence, it remains an important open problem to design better online algorithms for less restrictive CTRs.

The second line of improvements comes from combining the online primal-dual framework and a more expressive linear program for the problem. This new analysis method gives a better understanding of classical algorithms BALANCE [Huang and Zhang 2020] and RANKING [Huang et al. 2023] in the presence of stochastic rewards.

Researchers have also tried to gain new insight by considering a weaker clairvoyant benchmark. [Goyal and Udwani 2023] showed that against the weaker benchmark, RANKING achieves the optimal $1 - 1/e$ competitive ratio for uniform CTRs. They also analyzed BALANCE for small CTRs, and obtained a ratio better than the aforementioned state-of-the-art against the offline optimum benchmark. The latter result was later improved by [Huang et al. 2023] to 0.611.

Last but not least, stochastic rewards are closely related to budget-oblivious algorithms for AdWords, i.e., algorithms which do not know an agent's budget until the moment it is depleted. By a reduction by [Mehta et al. 2014], a competitive online algorithm for the latter model would yield the same competitive ratio in the former model (but not vice versa). Again, the greedy algorithm is a $1/2$ -competitive budget-oblivious algorithm for this problem. The survey by [Mehta 2013] listed finding a better budget-oblivious algorithm as an open problem. [Vazirani 2023] suggested a variant of RANKING as a candidate algorithm. [Liang et al. 2023] showed that no variant of this algorithm is $(1 - 1/e)$ -competitive. Finally, [Udwani 2023] proved that the candidate is at least 0.508-competitive, and a variant of this algorithm is 0.522-competitive, both under the small-bids assumption, whereby the bids w_{ai} 's are small compared to the agent's budget B_a .

3. BEYOND ONLINE BIPARTITE MATCHING AND AD ALLOCATION

The preceding online matching models, largely motivated by online advertising, crucially rely on the assumption that one side of the bipartite graph is fixed and known upfront. This prevents the theory of online matching being applied to other modern applications, including ride-hailing, ride-sharing, rental services, etc.

In this section, we discuss generalizations of classic online bipartite matching. The first two generalizations are motivated by ride-hailing and ride-sharing, that allow all vertices to arrive online and allow general (non-bipartite) graphs. The third is somewhat theoretical in nature, but is the most general problem. The last generalization is motivated by rental services and freelance labor markets, and so captures the reusability of resources.

3.1 Fully Online Model: Vertices with Arrivals and Deadlines

In online ride-hailing platforms (e.g., Uber, Lyft, DiDi), ride requests are submitted to the platform in an online fashion and are active in the system for a few minutes. The platform assigns each request to a currently available taxi (or self-employed driver). Requests and taxis can be modeled as vertices in a bipartite

	Fractional Relaxation	Original Problem
Fully Online	0.6 [Tang and Zhang 2022]	0.569 [Huang et al. 2020]
General Vertex Arrival	0.526 [Wang and Wong 2015]	$1/2 + \Omega(1)$ [Gamlath et al. 2019]
Edge Arrival	$1/2$ [Gamlath et al. 2019]	$1/2$ [Gamlath et al. 2019]
Reusable Resources	$1 - 1/e$ [Goyal et al. 2021]	0.589 [DeLong et al. 2023]
	$1 - 1/e$ [Feng et al. 2021]	

Table II. State-of-the-art for online matching problems beyond bipartite matching & ad allocation.

graph with edges between compatible ride requests and taxis. This is an online bipartite matching problem but does not fit into the classic model, since all vertices (both the requests and the taxis) arrive online. Similarly, ride-sharing platforms, which match ride requests (pairing up riders) are naturally modeled as an online matching problem on *general (non-bipartite) graphs*.

[Huang et al. 2020] introduced the fully online matching model to capture the above scenarios, though the same model was studied earlier by [Blum et al. 2006] in the context of liquidity in clearing markets. Let $G = (V, E)$ be the underlying graph, initially *completely unknown*. Each time step is either an arrival or a deadline of a vertex. Upon the arrival of a vertex, its incident edges to their previously-arrived neighbors are revealed. A vertex can be matched at any point until its *deadline*, with this time revealed on its arrival. Naturally, we assume the deadline of a vertex is after its arrival, and all edges incident to a vertex are revealed before its deadline. This model generalizes the classic one-sided online bipartite matching model, where all offline vertices arrive at the beginning and have deadlines at the end, and every online vertex has its deadline right after its arrival.

For the fully online matching problem, [Huang et al. 2020] proved that RANKING achieves a tight $\Omega \approx 0.567$ (the unique solution to $\Omega \cdot e^\Omega = 1$) competitive ratio for bipartite graphs, and a competitive ratio of 0.521 for general graphs. For the fractional variant of the problem, [Huang et al. 2019] established a tight $2 - \sqrt{2} \approx 0.585$ competitive ratio of BALANCE. Remarkably, RANKING and BALANCE are known to be optimal in the classic model, but the claimed tightness here only applies to the two algorithms themselves. Indeed, [Huang et al. 2020] introduced the BALANCED-RANKING algorithm that achieves a competitive ratio of 0.569 for bipartite graphs, and the EAGER WATER-FILLING algorithm that achieves a competitive ratio of 0.592 for the fractional variant. The later result was further improved by [Tang and Zhang 2022] to 0.6. On the negative side, the state-of-the-art upper bound (i.e., hardness) is 0.613 [Huang et al. 2020; Eckl et al. 2021; Tang and Zhang 2022], separating the fully online model from the classic online bipartite matching model.

This setting is also known as the windowed online matching problem. [Ashlagi et al. 2023] assumed a first-in-first-out structure on the active windows (i.e., arrivals and deadlines) of vertices, and achieved a $1/4$ competitive ratio for edge-weighted graphs through a reduction to the Display Ads problem by suffering an extra factor of 2. Combined with the state-of-the-art algorithm for the Display Ads problem by [Blanc and Charikar 2021], their competitive ratio can be improved to 0.268.

3.2 General Vertex Arrival

Generalizing fully online matching is the online matching with general vertex arrivals problem, introduced by [Wang and Wong 2015]. Again, the input is a graph $G = (V, E)$, initially unknown, with vertices arriving online. Upon the arrival of a vertex v , its incident edges to its previously-arrived neighbors are revealed. The algorithm either matches v to an unmatched neighbor immediately or leaves v unmatched, possibly matching it to a later-arriving neighbor u upon u 's arrival. The inability to match vertices at any point before their departure (and lack of this information) makes this model more restrictive than the fully online model, and so algorithms for general vertex arrivals are also algorithms in the fully online model, with the same competitive ratio.

[Wang and Wong 2015] presented a fractional 0.526-competitive algorithm for the fractional version of the problem. [Gamlath et al. 2019] designed a rounding of Wang and Wong's fractional algorithm and established a $1/2 + \Omega(1)$ competitive ratio for the integral matching problem. This result stands as the only non-trivial integral algorithm so far. On the negative side, [Wang and Wong 2015; Buchbinder et al. 2019; Tang et al. 2022] established an upper bound of 0.583, separating the general vertex arrival model from the fully online model.

3.3 Edge Arrivals

Finally, we remark that the most general online matching setting is the edge arrival model. That is, edges of an underlying graph arrive in a sequence and the algorithm decides whether to select an edge immediately on its arrival. Here edges correspond to fleeting collaboration opportunities between agents. A competitive ratio of $1/2$ can be trivially achieved by a greedy algorithm that matches each edge on arrival if both its endpoints are free, and this is optimal for deterministic algorithms. Unfortunately, [Gamlath et al. 2019] proved that no online algorithm achieves a better than $1/2 + 1/2n$ competitive ratio, even for the fractional version of the problem. Positive results are known assuming structure, including low-degree graphs and trees [Buchbinder et al. 2019], batching [Lee and Singla 2017], random-order arrivals [Guruganesh and Singla 2017], or stochastic arrivals [Gravin et al. 2021] (see §4 for more on the latter models).

3.4 Reusable Resources

In sponsored search, the advertisers' budgets, viewed as resources, are non-reusable. In contrast, in such markets as cloud computing (e.g., AWS, Azure), short-term rentals (e.g., Airbnb), and freelancer labor (e.g., TaskRabbit), the allocated resources (be it compute, housing or labor) are reusable, and can be reallocated after being used.

The above motivates online bipartite matching with reusable resources, where after an offline vertex (a rental service) is matched, it becomes available again after d time steps, where d is a known parameter that corresponds to the usage duration of the vertices. The classic online bipartite matching problem is a special case of the reusable resources model when $d = \infty$.

This model was first introduced by [Gong et al. 2022] in a more general setting of online assortment optimization. [Goyal et al. 2021; Feng et al. 2021] generalized

the BALANCE algorithm and achieves an optimal $1 - 1/e$ competitive ratio for the fractional version of the problem.² For the integral version of the problem, [Delong et al. 2023] proposed the PERIODIC RERANKING algorithm (a variant of RANKING that reranks the offline vertices every d time steps), and show that it achieves a competitive ratio of 0.589, and an online correlated selection-based algorithm achieves a competitive ratio of 0.505. All these results extend to the vertex-weighted setting. We remark that the results of [Delong et al. 2023] heavily rely on the assumption that all vertices have identical usage durations d . The case of heterogeneous usage durations (i.e., each vertex v has an individual duration time d_v) remains open.

[Feng et al. 2022] further studied online assortment of reusable resources in the stochastic setting. Reusable resources due to additional production have also been considered in infinite-horizon *stochastic* settings [Aouad and Saritaç 2020; Collina et al. 2020; Kessel et al. 2022; Patel and Wajc 2024]. We discuss stochastic settings (without reusable resources) in more detail in the following section.

4. STOCHASTIC MODELS: SECRETARIES, PROPHETS, AND PHILOSOPHERS

The preceding sections focused on adversarial models, where both input graph and arrival order are chosen by an adversary. This modeling choice, while robust, is quite pessimistic, and naturally results in worse guarantees than possibly achievable for real-world applications of interest. A natural way to obtain improved provable guarantees is to either consider random arrival orders (but adversarial input), or to posit a stochastic generative model, possibly learnt from historical data. Such models hearken back to classic results in optimal stopping theory concerning online Bayesian selection problems.

In the most basic setting, a buyer has a single item to sell, and impatient buyers arrive one after another and make take-it-or-leave-it bids for this single item. The buyer must select which bid to accept, immediately and irrevocably when the bid is made. Under adversarial models, a buyer cannot be competitive with the hindsight-optimal solution. In contrast, if the bids arrive in random order (referred to as the *secretary problem*), then a competitive ratio of $1/e$ is optimal [Dynkin 1963]. Similarly, if the successive bids v_i are drawn independently from known distributions $\mathcal{D}_i$, then the optimal competitive ratio is $1/2$ [Krengel and Sucheston 1978], i.e., the buyer can guarantee an expected gain at least half of that obtained by a “prophet” who knows the realization of the randomness,

$$\mathbb{E}[\text{Gain}] \geq \frac{1}{2} \cdot \mathbb{E}[\max_i v_i].$$

Such guarantees contrasting with the offline optimal, or prophet, are referred to as *prophet inequalities*. One may also contrast with the (computationally-unbounded) optimal online algorithm for such problems, which for reasons elaborated below we refer to as *philosopher inequalities*.

These models have been generalized and extended significantly over the years. In this section, we focus on recent developments for generalizations of the above to bipartite matching markets, where the buyer wishes to sell multiple heterogeneous items, and each arriving buyer proposes a different bid for each item. Put other-

²Equivalently, they assumed that resources (offline vertices) have large capacities.

wise, we focus on online bipartite matching models. We note that the buyer and seller terminology are not accidental, and these models have tight connections to questions in *mechanism design*, which we also discuss in this section.

	impossibility	algorithmic
Secretary Matching	$1/e$ [Dynkin 1963]	$1/e$ [Kesselheim et al. 2013]
Prophet Matching	$1/2$ [Krengel and Sucheston 1978]	$1/2$ [Feldman et al. 2015]
Philosopher Matching	0.99999 [Papadimitriou et al. 2021]	0.652 [Naor et al. 2023]

Table III. State-of-the-art for online bipartite weighted matching in stochastic settings.

4.1 Secretary Problems

For edge-weighted online bipartite matching with online vertices arriving in random order, [Korula and Pál 2009] were the first to obtain a constant-competitive ratio, specifically a $1/8$ -competitive algorithm. This was later improved to the optimal $1/e$ ratio by [Kesselheim et al. 2013], generalizing the classic single-item result of [Dynkin 1963], which we recall is the special case of a single offline vertex. For k heterogeneous offline vertices, [Kleinberg 2005] showed $(1 - 1/\sqrt{k})$ -competitive and truthful mechanism. [Ezra et al. 2022a] study secretary matching in general graphs (with vertices arriving with edges to their previously-arrived neighbors, as in §3.1 and §3.2). They show that the optimal competitive ratio in this level of generality is $5/12$, notably strictly greater than achievable for bipartite graphs.³

The random-order model similarly allows for improved guarantees for the special cases of vertex-weighted and unweighted online bipartite matching. For unweighted matching, [Guruganesh and Singla 2017] show that for *edges* revealed in random order a better than $1/2$ competitive ratio is possible, notably beyond the worst-case optimal for adversarial arrivals [Gamlath et al. 2019]. Similarly, random-order vertex arrivals in bipartite graphs (with arrivals on only one side of the graph) allow one to surpass the worst-case optimal $1 - 1/e$: a generalization of RANKING achieves a competitive ratio of 0.662 [Huang et al. 2019; Jin and Williamson 2021]. This generalizes results of [Karande et al. 2011; Mahdian and Yan 2011], who showed that for the unweighted problem RANKING (unchanged) achieves competitiveness beyond $1 - 1/e$, with the best known bound being 0.696 [Mahdian and Yan 2011]. As noted by these last two works, these results for random-order arrivals imply the same competitive ratios for stochastic matching problems with unknown i.i.d. distributions over arrival types. This remains the best known result for unknown distributions. In the following sections, we discuss the types of guarantees achievable under *known* distributions.

³For some intuition as to why this is not a contradiction, note that in the star example (i.e., the single-item problem), the center of the star arrives after the highest-bidding neighbor with probability $1/2 > 1/e$, and so greedily matching the center when it arrives is $1/2$ -competitive, and the lower bound of [Dynkin 1963] for bipartite graphs does not apply.

4.2 Prophet Inequalities

The optimal competitive ratio of $1/2$ for the single-item problem due to [Krengel and Sucheston 1978] was also obtained several years later using a single-threshold (i.e., posted-price) algorithm by [Samuel-Cahn 1984]. This results in truthful welfare-approximating mechanisms for single-item auctions. This connection between (pricing-based) prophet inequalities and mechanism design was later elaborated upon by researchers at the intersection of Economics and Computation [Hajiahayhi et al. 2007; Chawla et al. 2010; Kleinberg and Weinberg 2019]. Interestingly, very recently [Banihashem et al. 2024] show that *any* guarantee achieved by an online Bayesian selection algorithm can be achieved by a (dynamic) posted-price policy, implying that studying the non-strategic setting results in truthful mechanisms which achieve the same approximation of the social welfare as the algorithm in strategic settings.

The connection between (combinatorial) prophet inequalities and mechanisms design continues to motivate a flurry of results on prophet inequalities for increasingly involved markets, with more and more sophisticated combinatorial constraints on the sets of buyers that may be serviced, or items sold. See the excellent surveys [Hill and Kertz 1992; Correa et al. 2018] and [Hartline 2012; Lucier 2017] for more on prophet inequalities and their connection to mechanism design, respectively. In what follows, we focus on prophet inequalities subject to matching constraints.

For unweighted online bipartite matching, [Feldman et al. 2009] were the first to show that stochastic inputs allow for competitive ratio beyond the worst-case optimal $1 - 1/e \approx 0.632$. Specifically, they show that if online vertices' neighborhoods are drawn i.i.d. from a single known distribution, then a competitive ratio of 0.67 is achievable. There has been a long line of work studying this question, most recently [Jaillet and Lu 2013; Brubach et al. 2021; Huang and Shu 2021; Huang et al. 2022; Tang et al. 2022], with the current best competitive ratios being 0.7299 and 0.716 assuming integral and arbitrary arrival rates [Brubach et al. 2021; Huang et al. 2022].⁴ For edge-weighted matching, a number of results were obtained under integral arrival rates [Haeupler et al. 2011; Brubach et al. 2021], with the best ratio standing at 0.704, while for arbitrary arrivals the ratio of $1 - 1/e$ was only recently beaten [Yan 2024; Qiu et al. 2023]. In contrast, by a work of [Manshadi et al. 2012], no competitive ratios greater than $1 - 1/e^2 \approx 0.864$ and 0.823 are possible in the same settings.

For unweighted and vertex-weighted bipartite matching under (much more general) time-varying independent distributions, [Tang et al. 2022] recently provided the first algorithm surpassing the competitive ratio of $1 - 1/e$, presenting a 0.666-competitive algorithm. For edge-weighted matching a competitive ratio of $1/2$ is best possible, as this generalizes the single-item problem of [Krengel and Sucheston 1978]. This ratio is known to be achievable via numerous approaches [Feldman et al. 2015; Dütting et al. 2020; Ezra et al. 2022b]. This ratio of $1/2$ is even achievable under vertex arrivals in general graphs [Ezra et al. 2022b], or correlated arrivals [Aouad and Ma 2023]. In contrast, the problem is strictly harder under edge arrivals, where the best known competitive ratio is in the range $[0.344, 0.4]$ [MacRury et al. 2023]

⁴The arrival rate is the expected number of arrivals of a particular online type.

and $[0.349, 3/7]$ for bipartite graphs [MacRury et al. 2023; Correa et al. 2023]. For unweighted matching a competitive ratio of 0.502 is possible [Gravin et al. 2021].

4.3 Philosopher Inequalities

While a competitive ratio of $1/2$ is worst-case optimal for online Bayesian selection subject to bipartite matching constraints, this is still a pessimistic worst-case guarantee, as the lower bounds focus on worst-case distributions. The optimal algorithms for distributions of interest may allow for better competitive ratios. This optimal algorithm, which is the solution of a Markov Decision Problem (MDP), is computable in polynomial space via standard techniques. As shown by [Papadimitriou et al. 2021], this is the right characterization, and even *approximating* the optimal policy beyond some 0.999 ratio is PSPACE-complete (i.e., is as hard as the hardest problems requiring polynomial space). Hence, under standard complexity-theoretic assumptions, this optimal policy is not computable in polynomial time. Put otherwise, it is likely computable only by a character with sufficient time to “think” (i.e., compute), therefore naturally referred to as a “philosopher”. This motivates the study of polynomial-time approximation of the optimal online algorithm, which, in analogy with prophet inequalities (approximation of the optimal offline algorithm), we term *philosopher inequalities*.

[Anari et al. 2019] were the first to consider the approximation of the optimal online algorithm for online Bayesian selection. They considered bounded-depth and production-constrained laminar matroids, for which they provided $(1 + \epsilon)$ -approximate philosopher inequalities for any constant $\epsilon > 0$. [Dütting et al. 2023] obtained the same bounds for random-order (secretary) philosopher inequalities for a single item. For online bipartite matching, by [Papadimitriou et al. 2021] such an approximation would result in surprising developments in complexity theory. On the positive side, a successive line of work [Papadimitriou et al. 2021; Saberi and Wajc 2021; Braverman et al. 2022; Naor et al. 2023] showed that (increasingly) better than $1/2$ -approximate philosopher inequalities are possible. The current best bound stands at 0.652 (notably, above the natural bound of $1 - 1/e$ for online matching algorithms). We note that the recent work of [Banihashem et al. 2024] also translates the above (polynomial-time) policies approximating the optimal policy into pricing-based (and hence truthful) mechanisms providing the same approximation of the optimal mechanism.

5. OVERARCHING TECHNIQUES

5.1 Primal-Dual Algorithms

The primal-dual method has found wide applications in the area of online algorithms. Refer to [Buchbinder and Naor 2009] for a comprehensive survey. For online matching and related problems, the primal-dual schema was first adapted by [Buchbinder et al. 2007] to analyze BALANCE for the AdWords problem. We illustrate this idea in the special case of (unweighted) online bipartite matching.

We start with an economic interpretation of BALANCE. Consider the offline vertices as divisible items, and the online vertices as 0-1 unit-demand buyers. At any moment, each offline vertex v *prices* itself at $g(x_v)$ per (fractional) unit based on the current water level x_v (matched fraction of v) and thus its neighbor receives a *utility*

of $1 - g(x_v)$ per unit of v (fractionally) assigned to it, where $g(\cdot)$ is an increasing function. (Note that the price of v increases over time.) Upon its arrival, an online vertex u continuously chooses the unmatched neighbors giving u the largest utility. Recall that the dual of the maximum matching problem is the minimum vertex cover problem. The economic interpretation suggests a natural way to set the dual variables: for each offline vertex, let its dual variable be the total collected price, and for each online vertex, let its dual variable be the utility. The primal-dual framework asserts that in order to establish a Γ competitive ratio, it suffices to prove that the total gain of each item-buyer pair is Γ . To illustrate this approach, we provide a formal yet brief analysis in Appendix A.

Almost all fractional online matching algorithms were analyzed within the primal-dual framework. This includes algorithms for AdWords [Buchbinder et al. 2007], Display Ads [Devanur et al. 2016], fully online matching [Huang et al. 2019; Huang et al. 2020; Tang and Zhang 2022], general vertex arrivals [Wang and Wong 2015], stochastic matching [Tang et al. 2022], etc.

The primal-dual method for fractional matching crucially requires the dual constraints to be satisfied *always*. In contrast, [Devanur et al. 2013] noticed that it suffices to have the dual constraints hold *in expectation* for randomized algorithms, and used this observation to provide a simplified competitive analysis of RANKING for online (vertex-weighted) bipartite matching. Their approach is now referred to as the *randomized* primal-dual schema.

Their proof relied on an intuitive economic interpretation of RANKING that is similar to the economic interpretation of BALANCE. Instead of maintaining a dynamic price that depends on the water level, each offline vertex sets a randomized fixed price (according to the random permutation generated by RANKING) at the beginning. Then on the arrival of each online vertex, it buys the cheapest remaining neighbor. Again, we split the gain of each matched edge between its two endpoints (i.e., set the corresponding dual variables), according to the price of the offline vertex and the utility of the online vertex. For our EC readers, refer to [Eden et al. 2021] for a proof that is written explicitly in the language of price and utility and avoids duality.

A remarkable property of the randomized primal-dual schema is its intrinsic robustness for vertex-weighted graphs for all variants of online bipartite matching. Indeed, this schema often (if not always) provides a “free lunch”, allowing one to extend a result on unweighted graphs to vertex-weighted graphs while preserving the same competitive ratio. E.g., [Devanur et al. 2013; Huang and Zhang 2020; Huang and Shu 2021; Huang et al. 2022; Tang et al. 2022]. Going beyond the online bipartite matching model, in the fully online matching model [Huang et al. 2020] further developed the randomized primal-dual schema, by introducing a novel charging mechanic that allows a vertex other than the two endpoints of a matched edge to share the gain. [Levin and Wajc 2021] further found an application of the randomized primal dual framework for submodular maximization.

5.2 Randomized Rounding and Contention Resolution Schemes

The relax-and-round framework considers fractional relaxations as guides for randomized algorithms’ probabilistic choices. This section discusses the prevalence of this approach for online matching problems.

For bipartite matching, the standard relaxation allows us to assign a fractional value $x_e \in [0, 1]$ to each edge so that any node v has at most one unit assigned to its edges, $\sum_{e \ni v} x_e \leq 1$. Intuitively, x_e can be thought of as the marginal matching probability of edge e by some randomized algorithm, and use these fractions to obtain randomized algorithms. Indeed, since every fractional bipartite matching is the convex combination of integral matchings $M_1, \dots, M_k$, this intuition can be made formal, by randomly picking one such matching M_i with probability equal to its coefficient in the convex combination. This results in each edge e being matched with probability x_e , and thus preserves any linear objectives, $\sum_e w_e \cdot x_e$. We refer to such rounding schemes matching each edge with probability x_e as *lossless* rounding schemes.

Perhaps surprisingly, and as hinted at by Table I, for *online* edge-weighted bipartite matching and ad allocation problems, there exists a gap between our understanding of fractional algorithms and indivisible randomized algorithms. Indeed, as pointed out by [Devanur et al. 2013, Footnote 3], the above-mentioned integrality does not carry over to the online setting: for every randomized algorithm, there exist graphs on which the (optimally) $(1 - 1/e)$ -competitive fractional algorithm BALANCE achieves value $8/7$ times higher than any randomized algorithms. Therefore, rounding fractional algorithms seems to require losing a large multiplicative factor (in the worst case). At face value, this large gap seems to rule out the use of the relax-and-round approach to obtain good randomized (integral) online matching algorithms.

Despite the above, a large number of results in online matching in recent years are obtained by (or can be interpreted as employing) online randomized rounding of fractional solutions, often obtained using the primal-dual schema, §5.1. See [the FOCS23 workshop on the topic](#). There are three flavors of results in this vein.

5.2.1 Lossless rounding. The Online Correlated Selection (OCS) technique, elaborated upon in §5.4 [Fahrback et al. 2022; Gao et al. 2021; Blanc and Charikar 2021] can be seen as losslessly rounding *structured* fractional bipartite matchings in online settings. [Buchbinder et al. 2023] were the first to explicitly ask what structure is *necessary* for lossless online rounding of bipartite matching, i.e., allowing one to match each edge with probability *exactly* x_e . They considered other structured fractional online matching algorithms and provided lossless online rounding schemes for these, which they used to obtain generalizations of OCS and sharp randomness thresholds for beating deterministic algorithms for online bipartite matching. Similarly, “spread out” fractional matchings, e.g., ones assigning value $1/\Delta$ in graphs of maximum degree Δ , can be rounded *nearly* losslessly, i.e., one can match each edge e with probability $x_e \cdot (1 - \epsilon)$, and this is key to numerous results for online edge coloring, e.g., [Cohen and Wajc 2018; Wajc 2020; Blikstad et al. 2024].

5.2.2 Approximate Rounding. In the other extreme, [Naor et al. 2023] ask how well *arbitrary* fractional matchings $\vec{x}$ can be rounded online, and provide approximate rounding schemes that match each edge e with probability $0.652 \cdot x_e$, notably breaking the barrier of $1 - 1/e$ for this problem. They then use this scheme to obtain improved results for online edge coloring of multigraphs and philosopher inequalities, among others. Some prior results for philosopher inequalities [Papadim-

itriou et al. 2021; Saberi and Wajc 2021] are also obtained by such approximate online rounding schemes applied to LP relaxations incorporating constraints only applicable to online algorithms [Torricco and Toriello 2022; Buchbinder et al. 2014]. Similarly, the multiway OCS of [Gao et al. 2021] can be seen (and used) as an approximate rounding scheme that provides guarantees per offline vertex, as opposed to per edge, lending itself to results for vertex-weighted matching. The result of [Gamlath et al. 2019] for online matching under general vertex arrivals likewise follows a lossy rounding approach (applied to the fractional algorithm of [Wang and Wong 2015]), though here the approximation guarantees are more global than per-vertex or per-edge. More approximate rounding schemes are obtained by Online Contention Resolution Schemes (OCRS), whose guarantees are weak in the context of adversarial settings, but are central to prophet inequalities, as we now discuss.

5.3 Online Contention Resolution Schemes

Contention resolution schemes (CRS) have their origins in the (offline) submodular optimization literature [Chekuri et al. 2014], and follow a natural rounding approach: *activate* each element independently with probability x_e , and then select a high-valued feasible subset (in our case, a matching) among the active elements. This is obtained by guaranteeing each active element be selected with as high a probability as possible. This probability $\Pr[e \text{ selected} \mid e \text{ active}]$ is referred to as the *balance ratio* of the CRS.

The above approach can be generalized to online settings [Feldman et al. 2016], where inclusion of active element (in our case, edge) e must be made immediately upon its activation. By an approach due to [Yan 2011], using an appropriate convex *ex-ante relaxation*, an OCRS with balance ratio c provides c -approximate prophet inequalities for the same setting, by considering an element (buyer) active if their bid is in their x_e -th percentile. As [Lee and Singla 2018] show, the opposite is true: prophet inequalities that are c -competitive with respect to this relaxation yield OCRS that are c -balanced. By preceding discussions on pricing-based prophet inequalities being derivable from arbitrary prophet inequalities, we find that OCRS yield welfare-approximating truthful mechanisms [Banihashem et al. 2024]. Generally, OCRS have found widespread applications since their introduction. See [Patel and Wajc 2024] for a discussion.

When specified to matchings, for batched OCRS (capturing vertex arrivals), balance ratios of $1/2$ and $(1 + 1/e^2)/2 \approx 0.567$ are optimal for adversarial vertex arrivals in general graphs [Ezra et al. 2022b] and random-order vertex arrivals in bipartite graphs [MacRury and Ma 2024], respectively. For other settings, despite much progress in recent years, the optimal balance ratio is still unknown for adversarial edge arrivals [Gravin and Wang 2019; Correa et al. 2023; MacRury et al. 2023], random-order edge arrivals [Brubach et al. 2021; Pollner et al. 2022; MacRury et al. 2023] and random-order vertex arrivals [Fu et al. 2021; MacRury and Ma 2024].

5.4 Online Rounding: Online Correlated Selection

In this section we elaborate on one particular general rounding scheme for online bipartite matching and its generalizations, termed Online Correlated Selection (OCS). In each round, the OCS observes the online item *and a fractional allocation of this item* to the offline agents. It must then allocate this item in whole to one

of the agents. We measure the quality of an OCS by how an agent’s value for the allocated subset of items depends on the fractional allocation to the agent.

Consider unweighted matching as a running example. We expect that an offline agent gets matched with a higher chance as the total fractional allocation to the agent increases. Consider the following natural baseline algorithm: match each item i to an unmatched agent a interested in the item with probability proportional to x_{ai} , the fractional allocation of item i to agent a . Let x_a denote the total fractional allocation to agent a . This rounding algorithm guarantees that agent a is matched with probability at least $1 - e^{-x_a}$, and this bound is tight (c.f., [Gao et al. 2021]). Unfortunately, combining the Balance algorithm and this baseline bound only gives the trivial 0.5 competitive ratio. See [Gao et al. 2021] for the state-of-the-art techniques for designing OCS for unweighted and vertex-weighted matching, which improved the above bound for the probability of matching agent a , and the resulting competitive ratios for these two problems. [Hosseini et al. 2023] further applied the OCS technique to an unweighted online matching problem with class fairness as its main objective.

The definitions of OCS for display ads and AdWords are too technical to be covered concisely in this brief survey. We refer readers to [Fahrback et al. 2022] for the original definition of OCS for display ads and a proof of concept that non-trivial OCS exists, and [Blanc and Charikar 2021] for the best existing result along this line. See [Huang et al. 2020] for the definition of OCS for AdWords; the technique is similar to the proof of concept for display ads by [Fahrback et al. 2022].

Although the concept of OCS was first introduced for problems in non-stochastic models, the technique has found applications in stochastic models as well. In hindsight, this is perhaps unsurprising. The general recipe for online matching algorithms in stochastic models is to first solve an LP relaxation and then make online matching decisions taking the LP solution as a guide. By design, an OCS can treat the LP solution as a fractional allocation and convert it into online matching decisions. [Tang et al. 2022] used this approach to obtain the first non-trivial algorithm for the non-IID model of unweighted and vertex-weighted online stochastic matching. [Huang et al. 2022] further gave an OCS tailored for the stochastic model to get the best competitive ratios to date for IID unweighted and vertex-weighted online stochastic matching.

Acknowledgements

Zhiyi Huang is supported by an NSFC grant 6212290003. Zhihao Gavin Tang is supported by NSFC grant 61932002. David Wajc is supported in part by a Taub Family Foundation “Leader in Science and Technology” fellowship. We thank Niv Buchbinder for his short analysis of the BALANCE algorithm (in the appendix), and Sam Taggart for valuable feedback.

A. ADDENDUM: A TEACHABLE MOMENT

Following a quote attributed to Feynman, namely “If you want to master something, teach it,” we present short and self-contained (and in our opinion, quite teachable) proofs of two basic results in online bipartite matching: a competitive ratio of $1 - 1/e$ for BALANCE extended to vertex-weighted matching [Buchbinder et al. 2007], and an even shorter proof that no fractional algorithm can do better.

Recall that for the online vertex-weighted bipartite matching problem, each offline vertex i and its positive weight $w_i \geq 0$ are known up front. At each time t , online vertex t arrives, together with its edges $(i, t) \in E$ to its neighbors $i \in N(t)$, and we must decide to what extent $x_{i,t}$ to assign t to its neighbor i , from which the algorithm accrues a value of $w_i \cdot x_{i,t}$. Both offline vertices and online vertices must be assigned to a total extent of at most one.

A linear programming (LP) relaxation of the problem (allowing us to match each edge (i, t) to an extent of $x_{i,t}$), together with this LP's dual, are as follows.

$$\begin{aligned}
 \text{(P)} \quad & \max \sum_{(i,t) \in E} w_i \cdot x_{i,t} & \text{(D)} \quad & \min \sum_i y_i + \sum_t z_t \\
 \text{s.t.} \quad & \sum_t x_{i,t} \leq 1 \quad \forall i & \text{s.t.} \quad & y_i + z_t \geq w_i \quad \forall (i,t) \in E \\
 & \sum_i x_{i,t} \leq 1 \quad \forall t & & y, z \geq 0 \\
 & x \geq 0
 \end{aligned}$$

A.1 Short analysis of BALANCE

The BALANCE algorithm: Initialize zero primal and dual solutions, $\vec{x}$ and $\vec{y}, \vec{z}$. Let $g(x) := \frac{e^x - 1}{e - 1}$. For every online vertex t on arrival, letting $x_i := \sum_{t' < t} x_{i,t'}$ be the fractional degree of neighboring vertex $i \in N(t)$ before arrival of t , we increase $x_{i,t}$ for all $i \in A$, where

$$A := \arg \max_{i \in N(t)} \{w_i \cdot (1 - g(x_i + x_{i,t}))\},$$

so that this set A grows monotonically, until $\sum_i x_{i,t} = 1$ or $\min_{i \in N(t)} (x_i + x_{i,t}) = 1$.

Note that for the unweighted case ($w_i = 1$ for all i), this algorithm is precisely the WATER LEVEL algorithm described in §2. The intuition behind this algorithm is clear: since each offline vertex is equally likely to have no future neighbors, we wish to maximize the value assigned to edges of the least (fractionally) matched offline vertex, in case it (and it alone) has no future edges. That generalizing this approach allows to get a competitive ratio better than $1/2$ (also for vertex-weighted matching) is perhaps less immediate. We present a proof of this fact using LP duality, and specifically dual fitting.

Dual fitting: For our analysis (only), for each offline vertex i and online vertex t , we set dual values

$$\begin{aligned}
 y_i &\leftarrow w_i \cdot g\left(\sum_{t'} x_{i,t'}\right), \\
 z_t &\leftarrow \max_{i \in N(t)} \left(w_i \cdot \left(1 - g\left(\sum_{t' \leq t} x_{i,t'}\right) \right) \right).
 \end{aligned}$$

The first step of any primal-dual-based proof involves showing that the constructed dual is feasible, and hence its cost upper bounds the maximum gain (in hindsight).

LEMMA A.1. Vectors $\vec{y}, \vec{z}$ are dual feasible, i.e., they are positive and $y_i + z_t \geq w_i$ for all edges $(i, t) \in E$. Consequently, $\sum_i y_i + \sum_t z_t \geq OPT$, where OPT is the weight of a maximum vertex-weighted matching.

PROOF. As $g : [0, 1] \rightarrow [0, 1]$, we have that $y_i, z_t \geq 0$. On the other hand, by monotonicity of g (and positivity of $\vec{x}$), for every edge $(i, t) \in E$ we have that $y_i = w_i \cdot g(\sum_{t'} x_{i,t'}) \geq w_i \cdot g(\sum_{t' \leq t} x_{i,t'})$, and so

$$y_i + z_t \geq w_i \cdot g\left(\sum_{t' \leq t} x_{i,t'}\right) + w_i \cdot \left(1 - g\left(\sum_{t' \leq t} x_{i,t'}\right)\right) = w_i.$$

The lower bound $\sum_i y_i + \sum_t z_t \geq OPT$ then follows from weak LP duality, together with the primal LP being a fractional relaxation of the problem. $\square$

The second step in a primal-dual proof involves bounding the ratio of the primal and dual solutions' gain/cost. For online algorithms this typically boils down to bounding the ratio of the change in these values per cost, as in the following.

LEMMA A.2. For each time t , the increase in primal value, $\Delta P_t := \sum_i w_i \cdot x_{i,t}$, and the increase in the dual cost, $\Delta D_t := z_t + \sum_i (w_i \cdot g(x_i + x_{i,t}) - w_i \cdot g(x_i))$, satisfy

$$\Delta P_t / \Delta D_t \geq 1 - 1/e.$$

PROOF. The increase in dual cost satisfies

$$\begin{aligned} \Delta D_t &= \sum_i w_i \cdot (g(x_i + x_{i,t}) - g(x_i)) + z_t \\ &\leq \sum_i w_i \cdot x_{i,t} \cdot \left(g(x_i + x_{i,t}) + \frac{1}{e-1}\right) + z_t \\ &= \sum_i w_i \cdot x_{i,t} \cdot \left(g(x_i + x_{i,t}) + \frac{1}{e-1}\right) + \max_i w_i \cdot (1 - g(x_i + x_{i,t})) \\ &= \sum_i w_i \cdot x_{i,t} \cdot \left(g(x_i + x_{i,t}) + \frac{1}{e-1}\right) + \sum_i x_{i,t} \cdot \max_j w_j \cdot (1 - g(x_j + x_{j,t})) \\ &= \sum_i w_i \cdot x_{i,t} \cdot \left(g(x_i + x_{i,t}) + \frac{1}{e-1}\right) + \sum_i x_{i,t} \cdot w_i \cdot (1 - g(x_i + x_{i,t})) \\ &= \sum_i w_i \cdot x_{i,t} \cdot \left(1 + \frac{1}{e-1}\right) \\ &= \Delta P_t \cdot \frac{e}{e-1}. \end{aligned}$$

Above, the single inequality relied on the definition of $g(x) = \frac{e^x - 1}{e - 1}$ implying that $g'(x) = g(x) + \frac{1}{e-1}$ is monotone increasing in x . The second and third equalities follow by definition of z_t and either $\sum_i x_{i,t} = 1$ or $\min_{i \in N(t)} (x_i + x_{i,t}) = 1$, where the latter implies that $g(x_i + x_{i,t}) = 1$ for all $i \in N(t)$. The fourth equality follows from $x_{i,t} = 0$ unless $i \in \arg \max_{i \in N(t)} \{w_i \cdot (1 - g(x_i + x_{i,t}))\}$. $\square$

Combining both lemmas, the algorithm's competitive ratio follows.

THEOREM A.3. *Algorithm BALANCE is $(1 - 1/e)$ -competitive for vertex-weighted online bipartite matching.*

PROOF. Summing over all times t , we have by Lemmas A.2 and A.1 that

$$\sum_{i,t} w_i \cdot x_{i,t} = \sum_t \Delta P_t \geq (1 - 1/e) \cdot \sum_t \Delta D_t = \sum_i w_i \cdot g\left(\sum_{t'} x_{i,t'}\right) + \sum_t z_t \geq OPT,$$

where the last equality uses the telescoping sum

$$\sum_t \left(g\left(\sum_{t' \leq t} x_{i,t'}\right) - g\left(\sum_{t' < t} x_{i,t'}\right) \right) = g\left(\sum_{t'} x_{i,t'}\right) - g(0) = g\left(\sum_{t'} x_{i,t'}\right),$$

where the last equality follows from $g(0) = 0$. $\square$

A.2 Matching Impossibility Result

We now show that BALANCE is optimal, up to vanishingly small lower-order terms, even for the special case of *unweighted* online bipartite matching. The underlying input family is the same as that used by [Karp et al. 1990] to prove that randomized algorithms are at best $(1 - 1/e)$ -competitive. The following proof shows in a more direct manner that this ratio is best possible even for (potentially more powerful) fractional algorithms. The intuition behind the proof is precisely that guiding WATER LEVEL described above.

THEOREM A.4. *Any fractional online bipartite (unweighted) matching algorithm $\mathcal{A}$ has competitive ratio at most $1 - 1/e + o(1)$.*

PROOF. We consider inputs consisting of n offline and online vertices. The first online vertex neighbors all offline vertices; each following online vertex has the same neighborhood as its predecessor, barring one vertex. (Under appropriate labeling the “bipartite” adjacency matrix of this graph is upper triangular.) This graph clearly has a perfect matching, and therefore optimum value of n . To prove our theorem we show that for any fractional algorithm $\mathcal{A}$, a judicious choice of input forces $\mathcal{A}$ to achieve value at most $n(1 - 1/e) + 1$.

For every online vertex t , the neighbor i with no future edges is chosen to be a neighbor which was matched below the average, e.g., which minimizes $\sum_{t' \leq t} x_{i,t'}$. As the t -th online vertex to arrive neighbors $n - t + 1$ offline vertices, a proof by induction shows that offline vertices that do neighbor online vertex $t + 1$ are matched before time $t + 1$ to an average of at least $\sum_{t' \leq t} x_{i,t'} \geq \min\{1, \sum_{t'=1}^t \frac{1}{n-t'+1}\}$.⁵ Consequently, for $t = n(1 - 1/e) + 1$, every offline vertex i that neighbors $t + 1$ is fractionally matched to an extent of at least one, since

$$\sum_{t'=1}^{n(1-1/e)+1} \frac{1}{n-t'+1} = \sum_{i=n/e}^n \frac{1}{i} \geq \int_{n/e}^{n+1} \frac{1}{x} dx = \ln(n+1) - \ln(n/e) \geq 1.$$

⁵This holds only if the algorithm $\mathcal{A}$ is greedy, in the sense that it exhausts every online vertex t , i.e., sets $\sum_i x_{i,t} = 1$, when feasible. However, it is easy to modify any algorithm with a greedy algorithm which does no worse.

But since $\sum_{t'} x_{i,t'} \leq 1$, this implies that every offline vertex i neighboring online vertices after time t is already fully (fractionally) matched by time t , and so the algorithm gains no further profit from such i . Consequently, $\mathcal{A}$ achieves a value of at most $\sum_{i,t'} x_{i,t'} \leq n(1 - 1/e) + 1$, and so has competitive ratio at most $1 - 1/e + 1/n$. $\square$

See [Feige 2018] for tight bounds on the $o(1)$ term in the above upper bound.

REFERENCES

ReferencesRef

- AGGARWAL, G., GOEL, G., KARANDE, C., AND MEHTA, A. 2011. Online vertex-weighted bipartite matching and single-bid budgeted allocations. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1253–1264.
- ANARI, N., NIAZADEH, R., SABERI, A., AND SHAMELI, A. 2019. Nearly optimal pricing algorithms for production constrained and laminar bayesian selection. In *Proceedings of the 20th ACM Conference on Economics and Computation (EC)*. 91–92.
- AOUAD, A. AND MA, W. 2023. A nonparametric framework for online stochastic matching with correlated arrivals. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC)*. 114–114.
- AOUAD, A. AND SARITAÇ, Ö. 2020. Dynamic stochastic matching under limited time. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*. 789–790.
- ASHLAGI, I., BURQ, M., DUTTA, C., JAILLET, P., SABERI, A., AND SHOLLEY, C. 2023. Edge-weighted online windowed matching. *Mathematics of Operations Research* 48, 2, 999–1016.
- AZAR, Y. AND LITICHEVSKY, A. 2006. Maximizing throughput in multi-queue switches. *Algorithmica* 45, 1, 69–90.
- BANIHASHEM, K., HAJIAGHAYI, M., KOWALSKI, D. R., KRYSTA, P., AND OLKOWSKI, J. 2024. Power of posted-price mechanisms for prophet inequalities. In *Proceedings of the 35th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 4580–4604.
- BLANC, G. AND CHARIKAR, M. 2021. Multiway online correlated selection. In *Proceedings of the 62nd Symposium on Foundations of Computer Science (FOCS)*. 1277–1284.
- BLIKSTAD, J., SVENSSON, O., VINTAN, R., AND WAJC, D. 2024. Online edge coloring is (nearly) as easy as offline. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC)*. To Appear.
- BLUM, A., SANDHOLM, T., AND ZINKEVICH, M. 2006. Online algorithms for market clearing. *Journal of the ACM (JACM)* 53, 5, 845–879.
- BRAVERMAN, M., DERAKHSHAN, M., AND LOVETT, A. M. 2022. Max-weight online stochastic matching: Improved approximations against the online benchmark. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC)*. 967–985.
- BRUBACH, B., GRAMMEL, N., MA, W., AND SRINIVASAN, A. 2021. Improved guarantees for offline stochastic matching via new ordered contention resolution schemes. In *Proceedings of the 35th Annual Conference on Neural Information Processing Systems (NeurIPS)*.
- BUCHBINDER, N., JAIN, K., AND NAOR, J. S. 2007. Online primal-dual algorithms for maximizing ad-auctions revenue. In *Proceedings of the 15th Annual European Symposium on Algorithms (ESA)*. 253–264.
- BUCHBINDER, N., JAIN, K., AND SINGH, M. 2014. Secretary problems via linear programming. *Mathematics of Operations Research* 39, 1, 190–206.
- BUCHBINDER, N., NAOR, J., AND WAJC, D. 2023. Lossless online rounding for online bipartite matching (despite its impossibility). In *Proceedings of the 34th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 2030–2068.
- BUCHBINDER, N. AND NAOR, J. S. 2009. The design of competitive online algorithms via a primal-dual approach. *Foundations and Trends® in Theoretical Computer Science (TCS)* 3, 2–3, 93–263.
- BUCHBINDER, N., SEGEV, D., AND TKACH, Y. 2019. Online algorithms for maximum cardinality matching with edge arrivals. *Algorithmica* 81, 5, 1781–1799.

- CHAWLA, S., HARTLINE, J. D., MALEC, D. L., AND SIVAN, B. 2010. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the 42nd Annual ACM Symposium on Theory of Computing (STOC)*. 311–320.
- CHEKURI, C., VONDRÁK, J., AND ZENKLUSEN, R. 2014. Submodular function maximization via the multilinear relaxation and contention resolution schemes. *SIAM Journal on Computing (SICOMP)* 43, 6, 1831–1879.
- COHEN, I. R. AND WAJC, D. 2018. Randomized online matching in regular graphs. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 960–979.
- COLLINA, N., IMMORLICA, N., LEYTON-BROWN, K., LUCIER, B., AND NEWMAN, N. 2020. Dynamic weighted matching with heterogeneous arrival and departure rates. In *Proceedings of the 16th Conference on Web and Internet Economics (WINE)*. 17–30.
- CORREA, J., CRISTI, A., FIELBAUM, A., POLLNER, T., AND WEINBERG, S. M. 2023. Optimal item pricing in online combinatorial auctions. *Mathematical Programming*, 1–32.
- CORREA, J., FONCEA, P., HOEKSMAN, R., OOSTERWIJK, T., AND VREDEVELD, T. 2018. Recent developments in prophet inequalities. *SIGecom Exchanges* 17, 1, 61–70.
- DELONG, S., FARHADI, A., NIAZADEH, R., SIVAN, B., AND UDWANI, R. 2023. Online bipartite matching with reusable resources. *Mathematics of Operations Research*.
- DEVANUR, N. AND MEHTA, A. 2022. Online matching in advertisement auctions. In *Online and Matching-Based Market Design*. Cambridge University Press.
- DEVANUR, N. R., HUANG, Z., KORULA, N., MIRROKNI, V. S., AND YAN, Q. 2016. Whole-page optimization and submodular welfare maximization with online bidders. *ACM Trans. Economics and Comput.* 4, 3, 14:1–14:20.
- DEVANUR, N. R., JAIN, K., AND KLEINBERG, R. D. 2013. Randomized primal-dual analysis of ranking for online bipartite matching. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 101–107.
- DÜTTING, P., FELDMAN, M., KESSELHEIM, T., AND LUCIER, B. 2020. Prophet inequalities made easy: Stochastic optimization by pricing nonstochastic inputs. *SIAM Journal on Computing (SICOMP)* 49, 3, 540–582.
- DÜTTING, P., GERGATSOULI, E., REZVAN, R., TENG, Y., AND TSIGONIAS-DIMITRIADIS, A. 2023. Prophet secretary against the online optimal. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC)*. 561–581.
- DYNNIN, E. B. 1963. The optimum choice of the instant for stopping a markov process. *Soviet Mathematics* 4, 627–629.
- ECHENIQUE, F., IMMORLICA, N., AND VAZIRANI, V. V. 2023. *Online and matching-based market design*. Cambridge University Press.
- ECKL, A., KIRSCHBAUM, A., LEICHTER, M., AND SCHEWIOR, K. 2021. A stronger impossibility for fully online matching. *Operations Research Letters* 49, 5, 802–808.
- EDEN, A., FELDMAN, M., FIAT, A., AND SEGAL, K. 2021. An economics-based analysis of ranking for online bipartite matching. In *Proceedings of the 4th Symposium on Simplicity in Algorithms (SOSA)*. 107–110.
- EZRA, T., FELDMAN, M., GRAVIN, N., AND TANG, Z. G. 2022a. General graphs are easier than bipartite graphs: Tight bounds for secretary matching. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC)*. 1148–1177.
- EZRA, T., FELDMAN, M., GRAVIN, N., AND TANG, Z. G. 2022b. Prophet matching with general arrivals. *Mathematics of Operations Research* 47, 2, 878–898.
- FAHRBACH, M., HUANG, Z., TAO, R., AND ZADIMOGHADDAM, M. 2022. Edge-weighted online bipartite matching. *Journal of the ACM (JACM)* 69, 6, 1–35.
- FEIGE, U. 2018. Tighter bounds for online bipartite matching. *arXiv preprint arXiv:1812.11774*.
- FELDMAN, J., KORULA, N., MIRROKNI, V., MUTHUKRISHNAN, S., AND PÁL, M. 2009. Online ad assignment with free disposal. In *Proceedings of the 5th Conference on Web and Internet Economics (WINE)*. 374–385.
- FELDMAN, J., MEHTA, A., MIRROKNI, V., AND MUTHUKRISHNAN, S. 2009. Online stochastic matching: Beating $1-1/e$. In *Proceedings of the 50th Symposium on Foundations of Computer Science (FOCS)*. 117–126.

- FELDMAN, M., GRAVIN, N., AND LUCIER, B. 2015. Combinatorial auctions via posted prices. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 123–135.
- FELDMAN, M., SVENSSON, O., AND ZENKLUSEN, R. 2016. Online contention resolution schemes. In *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1014–1033.
- FENG, Y., NIAZADEH, R., AND SABERI, A. 2021. Robustness of online inventory balancing algorithm to inventory shocks. Available at SSRN: <https://ssrn.com/abstract=3795056>.
- FENG, Y., NIAZADEH, R., AND SABERI, A. 2022. Near-optimal bayesian online assortment of reusable resources. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC)*. 964–965.
- FU, H., TANG, Z. G., WU, H., WU, J., AND ZHANG, Q. 2021. Random order vertex arrival contention resolution schemes for matching, with applications. In *Proceedings of the 48th International Colloquium on Automata, Languages and Programming (ICALP)*.
- GAMLATH, B., KAPRALOV, M., MAGGIORI, A., SVENSSON, O., AND WAJC, D. 2019. Online matching with general arrivals. In *Proceedings of the 60th Symposium on Foundations of Computer Science (FOCS)*. 26–38.
- GAO, R., HE, Z., HUANG, Z., NIE, Z., YUAN, B., AND ZHONG, Y. 2021. Improved online correlated selection. In *Proceedings of the 62nd Symposium on Foundations of Computer Science (FOCS)*. 1265–1276.
- GONG, X., GOYAL, V., IYENGAR, G. N., SIMCHI-LEVI, D., UDWANI, R., AND WANG, S. 2022. Online assortment optimization with reusable resources. *Management Science* 68, 7, 4772–4785.
- GOYAL, V., IYENGAR, G., AND UDWANI, R. 2021. Asymptotically optimal competitive ratio for online allocation of reusable resources. In *Proceedings of the 17th Conference on Web and Internet Economics (WINE)*. 543.
- GOYAL, V. AND UDWANI, R. 2023. Online matching with stochastic rewards: optimal competitive ratio via path-based formulation. *Operations Research* 71, 2, 563–580.
- GRAVIN, N., TANG, Z. G., AND WANG, K. 2021. Online stochastic matching with edge arrivals. In *Proceedings of the 48th International Colloquium on Automata, Languages and Programming (ICALP)*. 74:1–74:20.
- GRAVIN, N. AND WANG, H. 2019. Prophet inequality for bipartite matching: Merits of being simple and non adaptive. In *Proceedings of the 20th ACM Conference on Economics and Computation (EC)*. 93–109.
- GURUGANESH, G. P. AND SINGLA, S. 2017. Online matroid intersection: Beating half for random arrival. In *Proceedings of the 19th Conference on Integer Programming and Combinatorial Optimization (IPCO)*. 241–253.
- HAEUPLER, B., MIRROKNI, V. S., AND ZADIMOGHADDAM, M. 2011. Online stochastic weighted matching: Improved approximation algorithms. In *Proceedings of the 7th Conference on Web and Internet Economics (WINE)*. 170–181.
- HAJIAGHAYI, M. T., KLEINBERG, R., AND SANDHOLM, T. 2007. Automated online mechanism design and prophet inequalities. In *Proceedings of the 22nd AAAI Conference on Artificial Intelligence (AAAI)*. 58–65.
- HARTLINE, J. D. 2012. Approximation in mechanism design. *American Economic Review* 102, 3, 330–36.
- HILL, T. P. AND KERTZ, R. P. 1992. A survey of prophet inequalities in optimal stopping theory. *Contemporary Mathematics* 125, 191–207.
- HOSSEINI, H., HUANG, Z., IGARASHI, A., AND SHAH, N. 2023. Class fairness in online matching. In *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 37. 5673–5680.
- HUANG, Z., JIANG, H., SHEN, A., SONG, J., WU, Z., AND ZHANG, Q. 2023. Online matching with stochastic rewards: Advanced analyses using configuration linear programs. In *Proceedings of the 20th Conference on Web and Internet Economics (WINE)*. 384–401.
- HUANG, Z., KANG, N., TANG, Z. G., WU, X., ZHANG, Y., AND ZHU, X. 2020. Fully online matching. *Journal of the ACM (JACM)* 67, 3, 1–25.

- HUANG, Z., PENG, B., TANG, Z. G., TAO, R., WU, X., AND ZHANG, Y. 2019. Tight competitive ratios of classic matching algorithms in the fully online model. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 2875–2886.
- HUANG, Z. AND SHU, X. 2021. Online stochastic matching, poisson arrivals, and the natural linear program. In *Proceedings of the 53rd Annual ACM Symposium on Theory of Computing (STOC)*. 682–693.
- HUANG, Z., SHU, X., AND YAN, S. 2022. The power of multiple choices in online stochastic matching. In *Proceedings of the 54th Annual ACM Symposium on Theory of Computing (STOC)*. 91–103.
- HUANG, Z., TANG, Z. G., WU, X., AND ZHANG, Y. 2019. Online vertex-weighted bipartite matching: Beating $1-1/e$ with random arrivals. *ACM Transactions on Algorithms (TALG)* 15, 3, 1–15.
- HUANG, Z., TANG, Z. G., WU, X., AND ZHANG, Y. 2020. Fully online matching ii: Beating ranking and water-filling. In *Proceedings of the 61st Symposium on Foundations of Computer Science (FOCS)*. 1380–1391.
- HUANG, Z. AND ZHANG, Q. 2020. Online primal dual meets online matching with stochastic rewards: Configuration lp to the rescue. In *Proceedings of the 52nd Annual ACM Symposium on Theory of Computing (STOC)*. 1153–1164.
- HUANG, Z., ZHANG, Q., AND ZHANG, Y. 2020. Adwords in a panorama. In *Proceedings of the 61st Symposium on Foundations of Computer Science (FOCS)*. 1416–1426.
- JAILLET, P. AND LU, X. 2013. Online stochastic matching: New algorithms with better bounds. *Mathematics of Operations Research*.
- JIN, B. AND WILLIAMSON, D. P. 2021. Improved analysis of ranking for online vertex-weighted bipartite matching in the random order model. In *Proceedings of the 17th Conference on Web and Internet Economics (WINE)*. 207–225.
- KALYANASUNDARAM, B. AND PRUHS, K. R. 2000. An optimal deterministic algorithm for online b -matching. *Theoretical Computer Science (TCS)* 233, 1, 319–325.
- KAPRALOV, M., POST, I., AND VONDRÁK, J. 2013. Online submodular welfare maximization: Greedy is optimal. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1216–1225.
- KARANDE, C., MEHTA, A., AND TRIPATHI, P. 2011. Online bipartite matching with unknown distributions. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing (STOC)*. 587–596.
- KARP, R. M., VAZIRANI, U. V., AND VAZIRANI, V. V. 1990. An optimal algorithm for on-line bipartite matching. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing (STOC)*. 352–358.
- KESSEL, K., SHAMELI, A., SABERI, A., AND WAJC, D. 2022. The stationary prophet inequality problem. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC)*. 243–244.
- KESSELHEIM, T., RADKE, K., TÖNNIS, A., AND VÖCKING, B. 2013. An optimal online algorithm for weighted bipartite matching and extensions to combinatorial auctions. In *Proceedings of the 21st Annual European Symposium on Algorithms (ESA)*. 589–600.
- KLEINBERG, R. AND WEINBERG, S. M. 2019. Matroid prophet inequalities and applications to multi-dimensional mechanism design. *Games and Economic Behavior* 113, 97–115.
- KLEINBERG, R. D. 2005. A multiple-choice secretary algorithm with applications to online auctions. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 630–631.
- KORULA, N. AND PÁL, M. 2009. Algorithms for secretary problems on graphs and hypergraphs. In *Proceedings of the 36th International Colloquium on Automata, Languages and Programming (ICALP)*. 508–520.
- KRENGEL, U. AND SUCHESTON, L. 1978. On semiamarts, amarts, and processes with finite value. *Probability on Banach spaces* 4, 197–266.

- LEE, E. AND SINGLA, S. 2017. Maximum matching in the online batch-arrival model. In *Proceedings of the 19th Conference on Integer Programming and Combinatorial Optimization (IPCO)*. 355–367.
- LEE, E. AND SINGLA, S. 2018. Optimal online contention resolution schemes via ex-ante prophet inequalities. In *Proceedings of the 26th Annual European Symposium on Algorithms (ESA)*. 57:1–57:14.
- LEVIN, R. AND WAJC, D. 2021. Streaming submodular matching meets the primal-dual method. In *Proceedings of the 32nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1914–1933.
- LIANG, J., TANG, Z. G., XU, Y. E., ZHANG, Y., AND ZHOU, R. 2023. On the perturbation function of ranking and balance for weighted online bipartite matching. In *Proceedings of the 31st Annual European Symposium on Algorithms (ESA)*. 80:1–80:15.
- LUCIER, B. 2017. An economic view of prophet inequalities. *ACM SIGecom Exchanges* 16, 1, 24–47.
- MACRURY, C. AND MA, W. 2024. Random-order contention resolution via continuous induction: Tightness for bipartite matching under vertex arrivals. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC)*. To Appear.
- MACRURY, C., MA, W., AND GRAMMEL, N. 2023. On (random-order) online contention resolution schemes for the matching polytope of (bipartite) graphs. In *Proceedings of the 34th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 1995–2014.
- MAHDIAN, M. AND YAN, Q. 2011. Online bipartite matching with random arrivals: an approach based on strongly factor-revealing lps. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing (STOC)*. 597–606.
- MANSHADI, V. H., GHARAN, S. O., AND SABERI, A. 2012. Online stochastic matching: Online actions based on offline statistics. *Mathematics of Operations Research* 37, 4, 559–573.
- MEHTA, A. 2013. Online matching and ad allocation. *Foundations and Trends® in Theoretical Computer Science* 8, 4, 265–368.
- MEHTA, A. AND PANIGRAHI, D. 2012. Online matching with stochastic rewards. In *Proceedings of the 53rd Symposium on Foundations of Computer Science (FOCS)*. 728–737.
- MEHTA, A., SABERI, A., VAZIRANI, U., AND VAZIRANI, V. 2007. Adwords and generalized online matching. *Journal of the ACM (JACM)* 54, 5, 22.
- MEHTA, A., WAGGONER, B., AND ZADIMOGHADDAM, M. 2014. Online stochastic matching with unequal probabilities. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1388–1404.
- NAOR, J. S., SRINIVASAN, A., AND WAJC, D. 2023. Online dependent rounding schemes. *arXiv preprint arXiv:2301.08680*.
- PAPADIMITRIOU, C., POLLNER, T., SABERI, A., AND WAJC, D. 2021. Online stochastic max-weight bipartite matching: Beyond prophet inequalities. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*. 763–764.
- PATEL, N. AND WAJC, D. 2024. Combinatorial stationary prophet inequalities. In *Proceedings of the 35th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 4605–4630.
- POLLNER, T., ROGHANI, M., SABERI, A., AND WAJC, D. 2022. Improved online contention resolution for matchings and applications to the gig economy. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC)*.
- QIU, G., FENG, Y., ZHOU, S., AND WU, X. 2023. Improved competitive ratio for edge-weighted online stochastic matching. In *Proceedings of the 20th International Conference on Web and Internet Economics (WINE)*. Springer, 527–544.
- SABERI, A. AND WAJC, D. 2021. The greedy algorithm is not optimal for online edge coloring. In *48th International Colloquium on Automata, Languages, and Programming (ICALP 2021)*. 109:1–109:18.
- SAMUEL-CAHN, E. 1984. Comparison of threshold stop rules and maximum for independent non-negative random variables. *the Annals of Probability* 12, 4, 1213–1216.

- SHIN, Y. AND AN, H. C. 2021. Making three out of two: three-way online correlated selection. In *Proceedings of the 32nd Annual International Symposium on Algorithms and Computation (ISAAC)*. 49.
- TANG, Z. G., WU, J., AND WU, H. 2022. (fractional) online stochastic matching via fine-grained offline statistics. In *Proceedings of the 54th Annual ACM Symposium on Theory of Computing (STOC)*. 77–90.
- TANG, Z. G. AND ZHANG, Y. 2022. Improved bounds for fractional online matching problems. *arXiv preprint arXiv:2202.02948*.
- TORRICO, A. AND TORIELLO, A. 2022. Dynamic relaxations for online bipartite matching. *INFORMS Journal on Computing* 34, 4, 1871–1884.
- UDWANI, R. 2023. Adwords with unknown budgets and beyond. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC)*. 1128–1128.
- VAZIRANI, V. V. 2023. Towards a practical, budget-oblivious algorithm for the adwords problem under small bids. In *Proceedings of the 43rd Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*.
- WAJC, D. 2020. Matching theory under uncertainty. Ph.D. thesis, Carnegie Mellon University.
- WANG, Y. AND WONG, S. C.-W. 2015. Two-sided online bipartite matching and vertex cover: Beating the greedy algorithm. In *Proceedings of the 42nd International Colloquium on Automata, Languages and Programming (ICALP)*. 1070–1081.
- YAN, Q. 2011. Mechanism design via correlation gap. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 710–719.
- YAN, S. 2024. Edge-weighted online stochastic matching: Beating. In *Proceedings of the 35th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 4631–4640.

Auto-bidding and Auctions in Online Advertising: A Survey

GAGAN AGGARWAL, ASHWINKUMAR BADANIDIYURU, SANTIAGO R. BALSEIRO, KSHIPRA BHAWALKAR, YUAN DENG, ZHE FENG, GAGAN GOEL, CHRISTOPHER LIAW, HAIHAO LU, MOHAMMAD MAHDIAN, JIEMING MAO, ARANYAK MEHTA, VAHAB MIRROKNI, RENATO PAES LEME, ANDRES PERLROTH, GEORGIOS PILIOURAS, JON SCHNEIDER, ARIEL SCHVARTZMAN, BALASUBRAMANIAN SIVAN, KELLY SPENDLOVE, YIFENG TENG, DI WANG, HANRUI ZHANG, MINGFEI ZHAO, WENNAN ZHU, and SONG ZUO

Google, Inc.

In this survey, we summarize recent developments in research fueled by the growing adoption of automated bidding strategies in online advertising. We explore the challenges and opportunities that have arisen as markets embrace this autobidding and cover a range of topics in this area, including bidding algorithms, equilibrium analysis and efficiency of common auction formats, and optimal auction design.

Categories and Subject Descriptors: J.4 [Social and Behavioral Sciences]: Economics

General Terms: Algorithms, Design, Economics, Theory

Additional Key Words and Phrases: Online advertising, autobidding, automation, auction design, price of anarchy

1. INTRODUCTION

Autobidding systems are taking on an increasingly large role in the online advertising ecosystem, with strong adoption by advertisers. Traditional bidding interfaces require advertisers to submit fine-grained bids, e.g., one bid per collection of keywords. With autobidding, an advertiser submits a high level goal and some high level constraints to the bidding platform. An *autobidding agent* for the advertiser then converts the goal and constraints into per-query bids at auction time, based on predictions of performance of each potential ad impression. Besides providing a much simplified interaction for the advertisers, autobidding also provides improved performance due to real-time optimization that takes predicted performance into account. Thus, it has already become a critical tool used by many advertisers.

There are several autobidding products in the advertising market. The oldest and most well-studied is that of budget optimization, in which the advertiser aims to maximize clicks from its ad campaign, and simply provides a daily budget and keywords to target. A more recent autobidding product is that of target-cost-per-acquisition (tCPA) in which the advertiser aims to maximize their post-click con-

Authors' addresses: {gagana, ashwinkumarbv, sbalseiro, kshipra, dengyuan, zhef, gagangoel, cvliaw, haihao, mahdian, maojm, aranyak, mirrokni, renatoppl, perlroth, gpil, jschnei, aschvartzman, balusivan, spendlove, yifengt, wadi, hanruiz, mingfei, wennanzhu, szuo}@google.com

versions (e.g., a sale or a sign-up), subject to an RoS (return-on-spend) constraint that the average cost per conversion is no more than a stated target. Target-return-on-ad-spend (tRoAS) generalizes tCPA to take into account the ultimate value of a conversion as well.

For each such product, the bidding agent automatically adjusts bids for its advertiser at auction time so as to maximize the performance for the campaign (under the given constraints), accounting for a dynamically changing environment such as query volume, query mix, or competition. We note that autobidding systems can be owned by the advertising platform as a service, or by third-parties. Given the importance of autobidding in the ad ecosystem, there has been considerable research in recent years to understand the fundamental properties, such as bidding algorithms, interaction with auction design, system equilibrium (i.e., the interaction across multiple autobidding agents for multiple advertisers), and mechanism design. This survey will attempt to cover a large portion of this growing and important literature.

2. PRELIMINARIES

In this section, we define the problem faced by the autobidding agents and the auctioneer in a set of unified notations. Consider the environment with n autobidding agents, indexed by i , and m auctions, indexed by j . The valuation of agent i winning auction j is $v_{ij} \in \mathbb{R}_+$. In general, the auctions are heterogeneous such that v_{ij} can vary with j . In the context of online advertising, the value v_{ij} may include predictions from machine learning models such as predicted click-through-rate and/or predicted conversion-rate, which can be different from auction to auction depending on the auction features made available to the prediction models.

Multi-slots. In the multi-slot environment, each auction can have up to $\ell \geq 1$ slots, indexed by k . The ℓ slots under each auction have decreasing importance to the agents because the ones in lower positions are less likely to attract the attention of the end user. This is modeled by a decaying factor $\beta_{jk} \in [0, 1]$ decreasing in k , such that the agent i winning the slot k of the auction j receives value $\beta_{jk} \cdot v_{ij}$. We assume that the decaying factor of the first slot is always normalized to 1, i.e., $\beta_{j1} := 1$. It is also without loss of generality to assume that each auction has the same number of slots, i.e., ℓ , because for any auction with fewer slots, one can always by introducing virtual slots with decay factor 0. To simplify the notations, in settings with only one slot, we will drop the decay factor β_{jk} .

Bidding and Auction. In each auction j , each agent submits a bid $b_{ij} \in \mathbb{R}_+$ and the auction takes the vector of bids $\mathbf{b}_j = (b_{1j}, \dots, b_{nj})$ as the input and determines the allocations $x_{ij}(\mathbf{b}_j) \in [0, 1]$ and payments $p_{ij}(\mathbf{b}_j) \in \mathbb{R}$ for each agent i . In the settings with multiple slots, the allocation becomes $x_{ijk}(\mathbf{b}_j) \in [0, 1]$, indicating the potentially randomized allocation of the slot k in auction j to agent i .

2.1 Bidding Agent's Problem

In the application of online advertising, the problem for the bidding agent is usually to maximize a given objective while subject to some constraints. There are two widely used objectives:

—*Utility-maximizing objective*: $\sum_{j \in [m]} x_{ij} \cdot v_{ij} - p_{ij}$;

—*Value-maximizing objective*: $\sum_{j \in [m]} x_{ij} \cdot v_{ij}$.

Value maximization focuses on maximizing clicks or conversions, regardless of cost, appealing to advertisers who prioritize these metrics. It indirectly considers payments through a constraint on payments or the return on spend, which we discuss next. Utility maximization, common in economics, maximizes the difference between value and payments, requiring values to be expressed in monetary terms, which can be difficult for advertisers. In certain settings, their hybrid version parameterized by $\lambda \in [0, 1]$ is also considered:

—*Hybrid objective*: $\sum_{j \in [m]} x_{ij} \cdot v_{ij} - \lambda \cdot p_{ij}$.

The constraints in practice can be designed to implement many different features, such as guaranteeing the number of wins, or limiting the maximum bids, etc. Out of which, the most commonly studied constraints are the *budget* constraint and the *return-on-spend (RoS)* constraint.

—*Budget constraint*: $\sum_{j \in [m]} p_{ij} \leq B_i$;

—*RoS constraint*: $\sum_{j \in [m]} x_{ij} \cdot v_{ij} \geq \tau_i \cdot \sum_{j \in [m]} p_{ij}$.

Budget constraints provide a natural way for advertisers to control their expenditures and are prevalent in advertising markets. We note that the RoS constraint has many equivalent forms, such as target CPA (cost-per-action) constraint, ROI (return-on-investment) constraint, IR (individual rationality) constraint with scaled values, etc. Throughout this survey, we will discuss all (equivalent) results in the language of the RoS constraints.

Taking advantage of the generality of the hybrid objective, we can formulate the bidding agent's problem as the following program:

$$\begin{aligned}
 \max \quad & \sum_{j \in [m]} x_{ij} \cdot v_{ij} - \lambda \cdot p_{ij} & (\text{BIDDING}) \\
 \text{s.t.} \quad & \sum_{j \in [m]} p_{ij} \leq B_i & (\text{BUDGET}) \\
 & \sum_{j \in [m]} x_{ij} \cdot v_{ij} \geq \tau_i \cdot \sum_{j \in [m]} p_{ij} & (\text{RoS})
 \end{aligned}$$

By picking different combinations of the parameters (λ, B_i, τ_i) , (BIDDING) can capture most of the settings that we are interested in.

— $\lambda = 0$: value-maximization;

— $\lambda = 1$: utility-maximization;

— $\tau_i = 0$: no RoS constraint;

— $B_i = +\infty$: no budget constraint.

Research Questions. From the bidders' perspective, the most important questions are:

- (1) *Optimal bidding* (Section 3.1): What is the optimal bidding strategy in a complete-information truthful auction?

- (2) *Online bidding for truthful auctions* (Section 3.2): How should agents bid in online truthful auctions when competition and valuations are uncertain?
- (3) *Online bidding for non-truthful auctions* (Section 3.3): How does non-truthfulness of the auction impacts agents' online learning strategies?

2.2 Auctioneer's Problem

Similar to the welfare (or liquid welfare in the presence of budgets) and revenue maximization in the canonical auction design settings, (liquid) welfare and revenue are also the most commonly concerned properties in autobidding environments. One major difference, when the bidding agent's objective is value-maximization with RoS constraint, is that the revenue and the (liquid) welfare are the same as long as the budget constraint or the RoS constraint binds. More specifically, define

—*Liquid welfare*: $\text{LWEL} = \sum_{i \in [n]} \min\{B_i, \sum_{j \in [m]} x_{ij} \cdot v_{ij} / \tau_i\}$;

—*Revenue*: $\text{REV} = \sum_{i \in [n]} \sum_{j \in [m]} p_{ij}$.

Liquid welfare is a measure of efficiency introduced by Dobzinski and Leme (2014), which quantifies the highest possible revenue that can be attained by a seller with full information on the bidders' information. In the literature, liquid welfare is used to measure efficiency instead of the usual social welfare because the latter cannot be well-approximated when bidders are constrained.

Observe that for each i ,

$$\sum_{j \in [m]} p_{ij} \leq \min\{B_i, \sum_{j \in [m]} x_{ij} \cdot v_{ij} / \tau_i\},$$

therefore $\text{REV} \leq \text{LWEL}$. The equality is reached when for all agent i , either (BUDGET) or (ROS) binds.

For value-maximization agents, i.e., $\lambda = 0$, under their optimal strategy (assuming others' fixed), it is often the case that either (BUDGET) or (ROS) binds, unless there is no chance for them to obtain higher values by increasing their spend. For this reason, (liquid) welfare receives significantly more attention in the literature, especially with value-maximization agents.

Research Questions. From the auctioneer's perspective, roughly three major categories of problems are concerned:

- (1) *Equilibrium* (Section 4.2): Does an equilibrium exist and, if so, it is unique and can it be efficiently computed? Do bidders converge to an equilibrium under different dynamics?
- (2) *Price of anarchy* (PoA, see Section 4.3 for formal definition): How are efficient equilibria in autobidding auctions? How auction design impacts the price of anarchy?
- (3) *Optimal auction design* (Section 5): How can we design auctions that improve the revenue or efficiency of the market?

2.3 Bayesian Auction Model

An alternative model widely adopted in the literature is a Bayesian model with a single auction in which bidder's values are drawn from a distribution $F \in \Delta(\mathbb{R}_+^n)$.

	m -auction model	Bayesian model
Hybrid objective	$\sum_{j \in [m]} x_{ij} \cdot v_{ij} - \lambda \cdot p_{ij}$	$\mathbb{E}[x_i \cdot v_i - \lambda \cdot p_i]$
Budget constraint	$\sum_{j \in [m]} p_{ij} \leq B_i$	$\mathbb{E}[p_i] \leq B_i/m$
RoS constraint	$\sum_{j \in [m]} x_{ij} \cdot v_{ij} \geq \tau_i \cdot \sum_{j \in [m]} p_{ij}$	$\mathbb{E}[x_i \cdot v_i] \geq \tau_i \cdot \mathbb{E}[p_i]$
Liquid welfare	$\sum_{i \in [n]} \min\{B_i, \sum_{j \in [m]} x_{ij} \cdot v_{ij}/\tau_i\}$	$\sum_{i \in [n]} \min\{B_i/m, \mathbb{E}[x_i \cdot v_i]/\tau_i\}$
Revenue	$\sum_{i \in [n]} \sum_{j \in [m]} p_{ij}$	$\sum_{i \in [n]} \mathbb{E}[p_i]$

Table I. Comparison between two models.

(For simplicity, we discuss the case of a single-slot auction.) The enumerative model defined above can be represented as a Bayesian model in which the random valuations $\tilde{\mathbf{v}} = (\tilde{v}_1, \dots, \tilde{v}_n)$ are drawn from a distribution F that takes value $(v_{1j}, \dots, v_{nj})$ with probability $1/m$ for each $j \in [m]$. Bayesian models are prominent in optimal auction design as they are useful to encode different informational assumptions and to represent settings with a continuum of valuations, which sometimes lead to more analytically tractable models.

Table I summarizes the key elements under the two different models. In the Bayesian model the budget and RoS constraints are naturally written in expectation over the realization of valuations and any randomness of the mechanism. This model can be interpreted as single-period problem with expected value constraints, instead of the more complex real-world scenario of multiple auctions with average constraints over time. This approach is more manageable and oftentimes allows for explicit solutions.

3. BIDDING ALGORITHMS

3.1 Optimal bidding for truthful auctions

In this section we present an LP based formulation of the autobidding problem from the view of an agent for one advertiser, slightly extending Aggarwal et al. (2019) to account for hybrid objectives.

Throughout this section, we will omit the subscript of i as we will take the perspective of a single bidding agent while assuming all other agents are fixed. Let p_j be the price of an ad for this advertiser for query j . Clearly, p_j depends on the bids of the other advertisers (who may also be using bidding agents), as well as on the pricing rule of the underlying auction. Suppose, for argument, we know the query sequence and the values of p_j in advance. Then we have the *selection problem* from Section 2.1: which queries would the advertiser like to buy so as to maximize their objective while staying within their constraints.

We rewrite the LP together with its Dual LP below, for a more generalized family of bidding problems with multiple constraints capturing different budget and RoS constraints $c \in \mathcal{C}$, parameterized by B^c , and w_j^c . The RoS constraint (like tCPA) is captured by setting the corresponding $B^c = 0$ and $w_j^c = \tau \cdot v_j$, and a budget constraint is captured by setting the corresponding B^c as the budget and $w_j^c = 0$. Finally, we remark that the extremal solutions of the LP are mostly integral if the query stream is large.

$$\begin{array}{l|l}
\begin{array}{l} \text{Maximize } \sum_j v_j x_j - \lambda p_j x_j \text{ s.t.} \\ \forall c \in \mathcal{C} : \sum_j p_j x_j \leq B^c + \sum_j w_j^c x_j \\ x_j \leq 1 \\ x_j \geq 0 \end{array} & \begin{array}{l} \text{Minimize } \sum_j \delta_j + \sum_c \alpha_c B^c \text{ s.t.} \\ \forall j : \delta_j \geq \sum_c \alpha_c (w_j^c - p_j) + v_j - \lambda p_j \\ \forall j : \delta_j \geq 0 \\ \forall c \in \mathcal{C} : \alpha_c \geq 0 \end{array}
\end{array} \quad \begin{array}{l} (1) \\ (4) \\ (2) \end{array} \quad (3)$$

In the dual problem, we denote by α_c the dual variable of the constraint $c \in \mathcal{C}$ and δ_j the dual variable of the constraint (2). We now leverage the LP formulation to come up with a bidding formula which can achieve the same optimal choice of queries as in the selection problem. The dual constraint (4) can be re-written as:

$$\forall j : \frac{\delta_j}{\lambda + \sum_c \alpha_c} \geq \left(\frac{v_j + \sum_c \alpha_c w_j^c}{\lambda + \sum_c \alpha_c} - p_j \right) \quad (5)$$

This directly gives us a bidding formula: Set the bid for query j ,

$$b_j := \frac{v_j + \sum_c \alpha_c w_j^c}{\lambda + \sum_c \alpha_c} \quad (6)$$

THEOREM 3.1. *Assuming that we have access to optimal values of the dual variables α_c , the bidding formula (6) results in an auction outcome identical to an optimal primal solution x_j , if the underlying auction is truthful.*

For a value maximizer ($\lambda = 0$) and simple RoS constraint with an additional a budget constraint (with corresponding optimal dual variables α_T and α_B , the bidding formula becomes:

$$b_j = \left(\frac{1 + \tau \cdot \alpha_T}{\alpha_T + \alpha_B} \right) v_j \quad (7)$$

For only an RoS constraint, the formula becomes $b_j = \left(\frac{1}{\alpha_T} + \tau \right) v_j$, essentially bidding proportionally to the value for an appropriate constant of proportionality. Finally, for a utility maximizer ($\lambda = 1$) with a budget constraint, we obtain the bidding formula

$$b_j = v_j / (1 + \alpha_B) \quad (8)$$

that was introduced by Balseiro et al. (2015).

The bid formula depends on the knowledge of the optimal duals α_c ; in practice these can be estimated via ML techniques from past data, and updated via control loops. We discuss the design of online algorithms in the next subsection.

3.2 Online learning for bidding in truthful auctions

There is a recent line of work studying the design of online learning algorithms under uncertainty. Most work in the literature consider a finite horizon model in which the bidder participates in m sequential auctions and constraints are enforced over time across auctions. For simplicity, we consider a single-slot auction. In this online model, when the j -th auction is announced, features are shared with the bidder and they estimate a value v_j for winning the auction. The value v_j is exogenously given and usually estimated using offline machine learning models. Valuation models tend to be more stable over time and can be trained across many

advertisers and long periods of time (McMahan et al., 2013; He et al., 2014; Zhou et al., 2018; Juan et al., 2016; Lu et al., 2017). The payment p_j is learned after the auction is cleared. While the value v_j is learned before bidding in the j -th auction, future values are not known in advance.

The online learning problem has been studied under different input models, stochastic and adversarial, and for truthful and non-truthful auctions. The case of non-truthful auctions is notoriously harder because the payment depends on the bid and the uniform bidding formula in (7) is not optimal—instead, the optimal bidding formula is a non-linear function of the value. We consider the case of non-truthful auctions in the next subsection.

Stochastic input. For truthful auctions, it is commonly assumed that the pairs $(v_j, p_j) \sim \mathcal{P}$ are independently and identically distributed (i.i.d.) from a joint distribution $\mathcal{P}$ that is unknown to the bidder. In other words, values and payments can be arbitrarily correlated for a given auction but independent across auctions. In this case, dual-based algorithms that update the dual variables α_B and α_T for the budget and RoS constraints, respectively, using first-order algorithms from online optimization has been shown to attain low regret relative to the offline optimization problem (BIDDING).

Balseiro and Gur (2019) consider the problem of a utility maximizer ($\lambda = 1$) with only a budget constraint ($\tau = 0$) and proposed *dual gradient descent*, a simple algorithm that adjusts the dual variable iteratively using online gradient descent. Denoting by $\alpha_{B,j}$ the value of the dual variable at the beginning of auction j , following (8), the bid is set to $b_j = v_j / (1 + \alpha_{B,j})$. Initially, the dual variable is set to $\alpha_{B,0} = 0$ and then it is updated as follows

$$\alpha_{B,j+1} = \max(\alpha_{B,j} + \eta \cdot (B/m - p_j), 0) ,$$

where η is a step-size that is usually chosen to be of order $\eta \approx m^{-1/2}$. Budget constraints are usually enforced strictly and the algorithm stops bidding whenever the budget is depleted. The term $B/m - p_j$ can be shown to be a subgradient of the j -th term of the dual objective (3) so the algorithm can be interpreted as performing gradient descent in the dual problem. The algorithm has an appealing self-correcting feature: it adjusts the dual variable to guarantee that the spend per auction is close to the average spend B/m . Balseiro and Gur (2019) show this algorithm obtains the following regret guarantee

$$\sup_{\mathcal{P}} \mathbb{E}_{(v_j, p_j) \sim \mathcal{P}} [\text{OPT} - \text{ALG}] = O(\sqrt{m}) ,$$

where OPT denotes the optimal objective value of the offline bidding problem (BIDDING). The result of Balseiro and Gur (2019) is proven under restrictive assumptions on the distributions $\mathcal{P}$, which were later relaxed by Balseiro et al. (2023).

Feng et al. (2023) provide a dual-based algorithm for a value maximizer with a budget and RoS constraint. They prove their algorithm attains regret $O(\sqrt{m})$ and incurs a violation of at most $O(\sqrt{m} \log m)$ of the RoS constraint. They also provide a more refined algorithm that satisfies both constraints strictly.

The near-optimal bidding algorithm of Feng et al. (2023) requires coordination between budget and RoS pacing systems to determine the bid. Balseiro et al. (2023) explore algorithms with different degrees of coordination between pacing systems.

In particular, they show that a fully-decoupled sequential algorithm could lead to poor performance and constraints violations, while a minimally-coupled algorithm that runs services independently can achieve similar performance to the optimal, fully-coupled algorithm.

Adversarial input. When values and payments are adversarially chosen, it is generally not possible to obtain fixed competitive ratios and, instead, one should settle for data-dependent competitive ratios. For value and utility maximizers with a budget constraint, Zhou et al. (2008) provide a dual-based algorithm that bids according to (8) and updates the dual variable based on the budget spent. Their algorithm attains a near-optimal competitive ratio of $1 + \log(U/L)$ where L and U are lower and upper bounds, respectively, on the value/utility to payment ratios. Their proof is established under the so-called “small bids assumption” that requires payments to be small relative to the budget. Later, Balseiro and Gur (2019) show that dual gradient descent obtains a competitive ratio of $\sup_j v_j/(B/m)$ for utility maximizers with a budget constraint. Their competitive ratio is also shown to be tight, which makes the algorithms of Zhou et al. (2008) and Balseiro and Gur (2019) not directly comparable.

3.3 Online learning for bidding in non-truthful auctions

When the auction is non-truthful, Theorem 3.1 does not hold and the optimal bid can be a complex function of values. A naïve approach is to reduce the problem to a contextual bandit with knapsacks in which the context is the value and each arm is a bid (Badanidiyuru et al., 2014). This approach requires discretization and results in a suboptimal regret bound of $O(m^{2/3})$ as it fails to exploit the structure of the problem. The most prominent non-truthful auction studied in the literature is the first-price auction in which the highest bidder wins and pays their value. First-price auctions have been recently adopted by many advertising platforms: Google switched to first-price auctions for its ad exchange in 2019 and Twitter switched in 2020 for mobile apps.¹

First-price auctions with budgets. Wang et al. (2023) focus on stochastic inputs where the values and the maximum competing bids are drawn from two independent distributions. For the full feedback model where the maximum competing bid is revealed after every auction, they provide a primal-dual algorithm that attains $\tilde{O}(\sqrt{m})$ regret. Using the graph-feedback and partial order properties in first-price auctions identified in Han et al. (2024), they also provide an algorithm with $\tilde{O}(\sqrt{m})$ regret in the one-sided feedback model where the bidders observe the maximum competing bid only if they lose the auction. Ai et al. (2022) study a model that additionally involves a discount factor in the objective function. They show $\tilde{O}(\sqrt{m})$ regret with full feedback and $\tilde{O}(m^{7/12})$ with one-sided feedback. Castiglioni et al. (2022) provide a general algorithm framework for the online knapsack problem with multiple resource constraints that can lead to low regret guarantees for stochastic inputs.

¹See <https://www.blog.google/products/admanager/rolling-out-first-price-auctions-google-ad-manager-partners> and <https://www.mopub.com/en/blog/first-price-auction>.

First-price auctions with budget and RoS constraints. Lucier et al. (2023) provide an algorithm with $\tilde{O}(m^{7/8})$ regret against the class of pacing multipliers, which bids uniformly and proportionally to the value. Remarkably, they prove a 2-approximation on the liquid welfare guarantee when all autobidders simultaneously adopt bid according to the proposed algorithm. Castiglioni et al. (2023) solve the general online learning problem under budget and RoS constraints. They endow standard primal-dual templates with weakly adaptive regret minimizers. Their framework applies to repeated first-price auctions where the set of possible valuations and bids are finite. Aggarwal et al. (2024) solve the problem of continuous valuations by designing an algorithm under full-information feedback, with $\tilde{O}(\sqrt{m})$ regret against the best possible Lipschitz function that maps values to bids. Their result builds on the primal-dual framework in Castiglioni et al. (2023) and is obtained by designing a dedicated tree-structured primal regret minimizer that achieves low interval regret. They also provide a lower bound of $\Omega(m^{2/3})$ regret with bandit feedback. Liang et al. (2023) design learning algorithms for an advertiser who repeatedly interacts with a platform when the selling mechanism/autobidding algorithm is a black box. They present a primal-dual algorithm for bandit feedback that attains good performance under different input models.

4. EQUILIBRIA AND PRICE OF ANARCHY

Within this section, most of the results only apply to the value-maximizing objective (i.e., $\lambda = 0$), and therefore, without explicit note, we assume $\lambda = 0$.

4.1 Solution Concepts

After defining the action space and the objective function for the agents, one natural question is to understand the game in terms of the properties of equilibria as well as other extended solution concepts. We summarize some solution concepts studied in the literature. We note that when either (BUDGET) or (RoS) is violated, the objective value for the agent is defined as $-\infty$. So in all the solution concepts below, (BUDGET) and (RoS) are forced to be satisfied by all agents.

- *Nash equilibrium* (NE): No agents can improve its objective value by changing to a different action $\mathbf{b}'_i$ within the set of (randomized) bids, i.e., $\mathbf{b}'_i \in \Delta(\mathbb{R}_+^m)$.
- *Pure Nash equilibrium* (PNE): No agent can improve its objective value by changing to a different action $\mathbf{b}'_i$ within the set of deterministic bids, i.e., $\mathbf{b}'_i \in \mathbb{R}_+^m$.
- *Undominated bids* (UDB) Balseiro et al. (2021a): No agent chooses a *dominated action*, i.e., $\mathbf{b}_i \in \text{UDB}_i$. An action $\mathbf{b}_i$ dominates another action $\mathbf{b}'_i$, if (i) For all possible bid profiles from others $\mathbf{b}_{-i} = (\mathbf{b}_1, \dots, \mathbf{b}_{i-1}, \mathbf{b}_{i+1}, \dots, \mathbf{b}_n)$, agent i 's objective value from $(\mathbf{b}_i, \mathbf{b}_{-i})$ is weakly higher than that from $(\mathbf{b}'_i, \mathbf{b}_{-i})$; (ii) And there exists one bid profile from others $\mathbf{b}''_{-i}$, such that agent i 's objective value from $(\mathbf{b}_i, \mathbf{b}''_{-i})$ is strictly higher than that from $(\mathbf{b}'_i, \mathbf{b}''_{-i})$. UDB_i is the set of actions of agent i that is not dominated by any other action.
- *PNE within UDB* (PNE + UDB): All agents choose *undominated* actions and no agent can improve its objective value by changing to a different *undominated* action, i.e., $\mathbf{b}'_i \in \mathbb{R}_+^m \cap \text{UDB}_i$.
- *PNE within Uniform Bidding* (PNE + UNI): All agents choose *uniform-bidding* actions and no agent can improve its objective value by changing to a different

uniform-bidding action $\mathbf{b}'_i \in \text{UNI}_i = \{\alpha \cdot \mathbf{v}_i : \alpha \in \mathbb{R}_+\}$.

- *Better-than-bidding-Values* (BTV) (Deng et al., 2022): No agent can improve its objective value by changing to the action of directly bidding its values, $\mathbf{b}'_i = \mathbf{v}_i$.
- *Autobidding Equilibrium* (ABE) (Li and Tang, 2024) (Second-price auction only): An ABE consists of the uniform-bidding actions for each agent and the allocation of the auctions such that: (i) Only agents with the highest bids can have non-zero allocation for the corresponding auctions, and all auctions are fully allocated; (ii) With the second price rule, the RoS constraints are respected and moreover, must bind unless the corresponding uniform-bidding factor reaches the upper limit.

We note that the relationship between these solution concepts could be complicated, and sometimes depends on the format of the auction. For example, PNE + UDB by definition is a subset of PNE, while PNE + UNI is not. In contrast, BTV is a necessary condition for a Nash equilibrium.

This section describes price of anarchy (PoA) results for a number of different solution concepts as described above. However, we note that our definition of PoA may not be “standard” since we may impose additional constraints that limit what equilibria are possible. Nonetheless, these results are useful from a practical point of view since the additional assumptions (such as UDB and UNI) are mild in practice.

4.2 Equilibrium existence and complexity

Aggarwal et al. (2019) shows the existence of PNE among autobidding agents in general multi-slot truthful auctions with two mild technical assumptions. One of them is, essentially, that there is no point mass in the value distributions. This assumption can be avoided by incorporating appropriate tie-breaking into the solution concept; this is an important advancement as a large body of the autobidding literature studies discrete instances, i.e., the values v_{ij} are deterministic and both n and m are finite. Conitzer et al. (2022) introduce this in the context of budget-constrained bidders and define the notion of a second-price pacing equilibrium (SPPE) – an SPPE is characterized by a vector of pacing multipliers as well as a fractional allocation of tied impressions that satisfies all the budget constraints. They prove the existence of SPPE for every pacing game, including those that are discrete and discontinuous, by constructing a sequence of smoothed games that converge to the (non-smoothed) pacing game. They show that a PNE exists for each of the smoothed games and the sequence of PNEs converges to an SPPE of the pacing game. Li and Tang (2024) extend this definition to RoS-constrained bidders, defining the term ABE, and also give a similar construction to prove existence of ABE for RoS-constrained bidders.

Another alternative to proving existence of an equilibrium is assuming a continuum of values so that ties are a zero-probability event. This approach has been successfully applied to utility-maximizers with budget constraints bidding in truthful auctions when values are independent or correlated but with positive densities and non-truthful standard auctions such as first-price auctions (see, e.g., Balseiro et al. 2015, 2021, 2023).

Chen et al. (2021) study the complexity of computing an equilibrium of the pacing game for budget-constrained bidders, and show that it is PPAD-hard to compute. Aggarwal et al. (2023) extend their result to show that it is PPAD-hard

to compute an ABE for RoS-constrained bidders. Li and Tang (2024) show that finding an approximate-ABE is also PPAD-hard.

4.3 Price of anarchy under different per-item auctions

Since Aggarwal et al. (2019), there are several lines of work that focus on the price of anarchy of canonical auction formats as well as their variants. Many of them consider the case with $\lambda = 0$ and are subject to the (RoS) constraint. Some of them consider the (BUDGET) constraint in addition.

To begin with, we first formally define the notion of *price of anarchy* (a.k.a. PoA) with respect to a solution concept. Let $\mathcal{I}$ denote the set of all autobidding environment instances, and $I \in \mathcal{I}$ denote one such instance. Let $\mathcal{A}$ denote an auction format, such as Vickrey-Clarke-Groves (VCG) auction, first price auction (FPA), etc. Let E be the solution concept and $E(\mathcal{A}, I)$ be the corresponding set of bid profiles under auction format $\mathcal{A}$ and instance I that satisfy the solution concept E .

Definition 4.1 Price of Anarchy. The price of anarchy of an auction $\mathcal{A}$ with respect to solution concept E is given by

$$\text{PoA}_E(\mathcal{A}) = \sup_{I \in \mathcal{I}, \mathbf{b} \in E(\mathcal{A}, I)} \frac{\max_{x^*} \text{LWEL}(x^*)}{\text{LWEL}(x^{\mathcal{A}}(\mathbf{b}))},$$

where $x^{\mathcal{A}}$ is the allocation function of the auction format $\mathcal{A}$, and the max is taken over all allocations x^* that are feasible with respect to the constraints. Note that here both x^* and LWEL depend on the instance I .

At a high-level, the price of anarchy tells us how much worse the worst-case equilibrium is from an optimal centralized allocation. The price of anarchy is always at least 1 and the closer it is to 1 the better.

In the above definition, “equilibrium” broadly refers to a solution concept detailed in Section 4.1. The specific concept used in each context will be clarified in the relevant section.

4.3.1 Basic auction formats: SPA, VCG, FPA, and GSP. The PoA result by Aggarwal et al. (2019) implies that in a second price auction, $\text{PoA}_{\text{PNE}+\text{UNI}}(\text{SPA}) \leq 2$. Furthermore, they show that this PoA bound is tight by showing that for $\varepsilon > 0$, there is an instance such that $\text{PoA}_{\text{PNE}+\text{UDB}+\text{UNI}}(\text{SPA}) \geq 2 - \varepsilon$. The original theorem is in fact for a much more general setting, where each bidding agent is subject to multiple general constraints that cover both (BUDGET) and (RoS) as special cases. The solution concept is then PNE plus the optimal bidding strategy based on the dual variables of the constraints, which degenerates to a uniform bidding strategy when one only has (BUDGET) and (RoS) as the constraints. A generalized version of liquid welfare is also used for defining the notion of PoA.

Deng et al. (2021) generalize the bound of 2 to multi-slot settings, and hence proves that $\text{PoA}_{\text{PNE}+\text{UNI}}(\text{VCG}) \leq 2$. This bound is tight as SPA can be considered as VCG on special single-slot instances. Hence $\text{PoA}_{\text{PNE}+\text{UNI}}(\text{VCG}) \geq \text{PoA}_{\text{PNE}+\text{UNI}}(\text{SPA}) = 2$. They also extend the results to VCG with certain additive boosts, which we will cover in Section 4.3.2.

Beyond truthful auctions, when there is only (RoS) but no (BUDGET), Liaw

et al. (2023) proves that $\text{PoA}_{\text{PNE}}(\text{FPA}) = 2$, and Deng et al. (2022) generalizes the result to randomized strategies, i.e., $2 \leq \text{PoA}_{\text{NE}}(\text{FPA}) \leq \text{PoA}_{\text{BTV}}(\text{FPA}) \leq 2$. When there are both agents with $\lambda = 0$ and $\lambda = 1$ in the environment, Deng et al. (2022) further shows that $\text{PoA}_{\text{NE}}(\text{FPA}) = 1 + \max_{t \in [0,1]} \frac{1-t}{1+t \ln t} \approx 2.188$. $\text{PoA}_{\text{NE}}(\text{FPA})$ can be improved with proper reserve prices, which we will cover in Section 4.3.2.

As a non-truthful generalization of SPA in multi-slot settings, GSP is another auction format commonly studied in the literature. Deng et al. (2024) proves an upper bound for $\text{PoA}_{\text{BTV}}(\text{GSP})$ depending on the decaying factors β_{jk} , which is unbounded in general after taking sup over all possible instances, i.e., $\text{PoA}_{\text{BTV}}(\text{GSP}) = \infty$. Deng et al. (2023) further show a fined-grained PoA with respect to the discount factors, i.e., the ratios of click probabilities between lower slots and the highest slot in each auction.

4.3.2 Basic auction formats with reserves and additive boosts. A line of research studies how PoA can be improved when the auctioneer has additional information about agent values and uses it as simple adjustments on basic auction formats.

Deng et al. (2021) shows that in position auctions with value maximizing agents with (ROS) constraints (and potentially also (BUDGET) constraints), for any constant $c > 0$, if the auctioneer sets an additive boost for each agent as c times the agent's value in VCG, $\text{PoA}_{\text{PNE+UNI}}$ is at most $(c+2)/(c+1)$. As $(c+2)/(c+1) < 2$ for any $c > 0$, this is a strict PoA improvement compared to VCG, and this PoA approaches 1 when c goes to infinity.

Balseiro et al. (2021a) improve upon Deng et al. (2021) in mainly three directions: (1) allowing auctioneer's additional information (a.k.a. ML advice) about agent values to be approximate, (2) studying reserves in addition to additive boosts, and (3) considering the mixed environment with general agent for $\lambda \in [0, 1]$. In particular, they show that if the auctioneer has a signal $\in [\gamma \cdot \text{value}, \text{value})$ for each agent, VCG with reserves or additive boosts gets PoA_{UDB} at most $2 - \gamma$, and VCG with both reserves and additive boosts gets PoA at most $2/(1 + \gamma)$. They also extend these results from VCG to GSP with slightly weaker guarantees. With such ML advice, Deng et al. (2022) also show this result can be extended to FPA with reserves setting to obtain $\text{PoA}_{\text{NE}} = \min_{t \in [0,1]} \frac{2-t-\gamma+(1-\gamma)t \ln t}{1+t \ln t - \gamma t(1+\ln t)}$.

In the presence of user costs, Deng et al. (2023) observe that the PoA of VCG can be arbitrarily bad (i.e., $\text{PoA}_{\text{PNE+UDB}}(\text{VCG}) = \infty$). They show that constant PoA can be restored by introducing either auction-dependent reserve prices or agent-dependent reserve prices.

4.3.3 Basic auction variants with randomization. The first paper to study randomized auctions in the autobidding setting is Mehta (2022). The mechanism (RAND) considered in this paper is defined by two parameters: a gap parameter α and a swap probability $p \leq 1/2$. If the gap between the highest bidder and the second highest bidder is at least α then the highest bidder wins. Otherwise, the highest bidder wins with probability p and with the remaining probability, the second highest bidder wins. The payment for the bidders is computed using Myerson's

payment rule.² They show that in the setting with two bidders, there is a choice of $\alpha \geq 1$ and p that gives a PoA of around 1.9. Note that even when there are only 2 bidders, the example in Aggarwal et al. (2019) shows that the PoA of SPA is 2.

A followup work by Liaw et al. (2023) considers mechanisms that are both randomized and non-truthful (i.e., the payment does not necessarily follow Myerson’s payment rule). More specifically, they consider a mechanism called randomized first-price auction (rFPA), which has an allocation function which is a generalization of RAND, but charges each winning bidder its bid. They show that, with an appropriate choice of their parameter α , this further improves the PoA to 1.8. We note that a key difference between Liaw et al. (2023) and Mehta (2022); Aggarwal et al. (2019) is that since the auction is no longer truthful, uniform bidding is not a best response and one has to analyze all possible bidding strategies.

We note that it is an open problem to design a randomized mechanism that has a PoA of *strictly* less than 2 for any fixed number n of bidders. A more difficult open problem is to exactly compute the PoA as a function of n . We remind the reader that these open problems are in the setting where the auction only receives bids and does not have any prior information on the values.

4.3.4 In the presence of budget constraint. Liaw et al. (2024) studies the auto-bidding settings with both (BUDGET) and (ROS). They first show that the gap between the optimal deterministic allocation and the optimal randomized allocation is n . Next, they define integral-PoA (I-PoA), which is the same as Definition 4.1, with an extra constraint that $x_{i,j}^* \in \{0, 1\}$. They show that the PoA of FPA is n , but it decreases to 2 under the mild assumption that for any bidder, their value for any query is at most their total budget. Interestingly, the I-PoA of FPA is 2 when there is only a single (ROS) constraint (Liaw et al., 2023) and when there are both (BUDGET) and (ROS) constraints. This means that the I-PoA does not get worse when the (BUDGET) budget constraint is added on top of the (ROS) constraint.

Uniform bidding is shown to be near optimal for bidders in truthful auctions Aggarwal et al. (2019), and achieves an optimal PoA of 1 for FPA with (ROS) constraints Deng et al. (2021). Liaw et al. (2024) shows that the I-PoA of FPA with uniform bidding is n , which is worse than the I-PoA = 2 for non-uniform bidding. The reason is that the bidders could be in a situation that they either win no query or would violate their budget by uniformly increasing bids for every query. However, uniform bidding improves the PoA for rFPA, because the bidders could increase bids smoothly to get more fraction value to avoid the bad cases in deterministic auctions. Finally, the authors propose a “quasi-proportional” FPA mechanism that achieves a PoA of 2 with both (BUDGET) and (ROS) constraints.

4.4 Bidding Dynamics

Even though equilibria are shown to exist under some mild conditions, it remains unclear whether the bidding agents will eventually converge to an equilibrium by following their bidding algorithms. The work of Paes Leme et al. (2024) shows that

²Fix a bidder i and let $x_i(b_i, b_{-i})$ be the allocation to bidder i when their bid is b_i and all other bids are b_{-i} . Myerson’s payment rule says that the payment to bidder i should be given by $b_i x_i(b_i, b_{-i}) - \int_0^{b_i} x(t, b_{-i}) dt$.

even with simple bidding algorithms, complex behavior can emerge in autobidding systems. For example, in one case of two bidders, there can be bi-stability (i.e., the existence of two stable equilibria), and the equilibrium to which the bidders converge depends upon the initial configuration of the multipliers. In the case of three bidders, they observe that there can be a stable periodic orbit, which implies that for some initial conditions the bidding system will never converge, even if an equilibrium does exist. Furthermore, they show that autobidding systems can simulate both linear dynamical systems as well as logical Boolean gates.

Liu and Shen (2023) study the optimal bidding strategy as the response to fixed strategies from competing agents in second price auctions. All agents have utility-maximization objectives under both budget and RoS constraints. When all agents adopt the proposed response strategy, they provide a sufficient condition such that the bidding dynamics converge to an equilibrium.

5. AUCTION DESIGN

Given the behavior of bidding agents defined by (BIDDING), a natural question is what are the efficient (optimal) auctions. From Section 4, we know that most of the commonly studied auction mechanisms are approximately efficient with constants at least 1.8. In this section, we introduce recent works on optimal auction design where the agents follow the optimization problem (BIDDING).

5.1 Bayesian auction design

In this subsection, we focus on the single Bayesian auction model introduced in Section 2.3, which is general enough to capture the discrete m -auction model.

In general, each agent i has three types of private information: (i) value v_i , (ii) budget B_i , and (iii) RoS target τ_i . Table II classifies the recent works based on whether each of these information is private or public as well as the choices of hybrid parameter λ . A distinctive feature of the auction design literature for autobidding auctions is the assumption that valuations are public instead of private as it is standard in the mechanism design literature. This assumption is predicated on the fact that advertisers increasingly rely on the machine learning algorithms that are developed by the advertising platforms to predict clicks and conversions.

value v_i	RoS target τ_i	budget B_i	λ	paper
private	public	$B_i = \infty$	$\lambda = 1$	Golrezaei et al. (2021)
public	private	$B_i = \infty$	$\lambda = 0$ or $\lambda = 1$	Balseiro et al. (2021b)
private	public	$B_i = \infty$	$\lambda = 0$ or $\lambda = 1$	Balseiro et al. (2021b)
private	public	$B_i = \infty$	$\lambda \in (0, 1)$	ex-post RoS Lv et al. (2023)
private	private	$B_i = \infty$	$\lambda = 0$	deterministic Balseiro et al. (2024)
private	public	public	$\lambda = 1$	Goel et al. (2014)
public	private	public	$\lambda = 0$	Balseiro et al. (2022)
public	private	private	$\lambda = 0$	Xing et al. (2023)

Table II. Relevant works by the information structure on values, RoS target, and budget, as well as the bidding agent objective type (parameterized by λ).

5.1.1 *RoS constraint only.* Golrezaei et al. (2021) consider the revenue-optimal auction design for utility-maximization agents ($\lambda = 1$) with ROI constraints, which can be equivalently modeled with RoS constraints. They find empirically, some buyers in the online ad market behave as if they are subject to such constraints. In the symmetric setting where agents have the same RoS target, they show that an optimal auction is one of the following depending on the RoS target: (i) second-price auction with the Myersonian reserve price, (ii) second-price auction with a reduced reserve price, (iii) second-price auction without reserve plus a participation subsidy. In the general asymmetric case, the optimal auction is more complex and can be interpreted in terms of modified virtual values.

Balseiro et al. (2021b) study the revenue-optimal mechanisms under different information structure on values and RoS targets for agents with either value-maximization objectives ($\lambda = 0$) or utility-maximization objectives ($\lambda = 1$). In the case of value-maximization ($\lambda = 0$), when either the values of agents are public information or the RoS targets of the agents are public information, they construct optimal mechanisms that achieve the first best (i.e., the optimal allocation when agent types are all public), which is not true in general when both values and RoS targets are private. In contrast, for the case of utility-maximization ($\lambda = 1$), when either the values of agents are public information or the RoS targets of the agents are public information, they construct the corresponding optimal mechanisms, while the first best cannot be achieved.

Lv et al. (2023) consider the revenue-optimal auction for bidding agents with intermediate objectives ($\lambda \in (0, 1)$) and require the RoS constraint to be satisfied ex-post instead of ex-ante, where the values of agents are private while the RoS targets of agents are public. They first provide a full characterization for dominant-strategy incentive compatibility: (i) monotone allocation rule and (ii) unique payment rule for any given monotone allocation. These can be seen as a generalization of Myerson's lemma (Myerson, 1981), while the unique payment rule follows a different relationship with the given allocation rule. They obtain the optimal auction for the single bidder case ($n = 1$) when a certain regularity condition is assumed (Decreasing Marginal Revenue).

Balseiro et al. (2024) prove that for the single value-maximization agent case ($n = 1$ and $\lambda = 0$), when both the values and RoS target are private information, the revenue-optimal mechanism with deterministic allocation can be implemented as a two-part tariff, i.e., a fixed price for buying the item and a fixed subsidy for not buying the item. An important implication from the structure of the optimal mechanism is that one does not need to screen the agent's RoS target.

5.1.2 *RoS and Budget constraints.* Goel et al. (2014) propose a generalized notion of admissible set that covers both the budget constraint and the RoS constraint. An admissible set can be modeled as $p_i \leq \alpha_i(x_i)$, where α_i is an increasing function. When it is a constant, it can capture the standard budget constraint, and when it is linear in x_i , it captures the RoS constraint. When it is the minimum of them, it captures both constraints at the same time. With the admissible set model, they consider the auction design with utility-maximization agents ($\lambda = 1$). In particular, they design a clinching auction (Ausubel, 2004) that is incentive compatible, individually rational and Pareto-efficient.

Balseiro et al. (2022) study the case with the budget constraint in addition to the RoS constraint. They consider the case for value-maximization agents ($\lambda = 0$) where the values and budgets of the agents are public information while the RoS targets are private. They obtain the revenue-optimal mechanism for $n = 1$ and $n = 2$ in general cases, and the optimal mechanism for $n \geq 3$ for special cases. Specifically, their optimal mechanism implements the efficient allocation according to RoS targets clipped up to thresholds depending on others' reports.

Xing et al. (2023) focus on the setting with value-maximization agents ($\lambda = 0$) where the values of agents are public information but both the RoS targets and budgets of the agents are private information. They provide the necessary and sufficient conditions for any allocation rule that can derive a truthful auction, and hence reduce the design space to allocation rules satisfying those conditions. Based on this characterization, they propose a family of simple truthful auctions. Although those auctions are not necessarily optimal, the characterization result is a non-trivial advancement towards this public value, private budget and RoS setting.

5.2 Auction design with ML advice

In online advertising, the auctioneer may have additional information about bidders' values via various machine learning technologies, i.e., ML advice. This additional information can be modeled as priors in a Bayesian setup as in Section 5.1.

Alternatively, Deng et al. (2021); Balseiro et al. (2021a); Deng et al. (2022) take a prior-free approach and model this ML advice as an approximate signal $\in [\gamma \cdot \text{value}, \text{value})$ for each bidder. They show using this ML advice as reserves or boosts in VCG and FPA can significantly improve welfare efficiency (see Section 4.3.2 for more details). With ML advice as reserves, Deng et al. (2022) demonstrate an individual welfare lower bound guarantee for this advertiser that increases in the advertiser's uniform bid multiplier, the quality of ML advice, and the relative market share of this advertiser compared to competitors. Together with results in Balseiro et al. (2021a), incorporating ML advice as personalized reserves achieves "best of both worlds" by simultaneously benefiting total and individual welfare.

5.3 Interdependent Auctions

Lu et al. (2023) consider a non-Bayesian model that is slightly different from our setting introduced in Section 2, where the allocation and payment in each single auction j depend on the bids $\{\mathbf{b}_i\}_{i=1}^n$ on all auctions. In other words, the allocation and payment of each single auction are no longer independent, instead, all the auctions are interdependent. They focus on constructing interdependent auctions with value-maximization agents ($\lambda = 0$) having a good competitive ratio compared against the offline optimal benchmark (i.e., no incentive constraints). They establish upper and/or lower bounds on the competitive ratios for several combinations across the information structure (fully private vs partially private), the demand type of agents (single-item, multi-item unit-demand, multi-item additive), and item divisibility.

5.4 Auctions with Alternative RoS Constraint

Wilkens et al. (2016, 2017) initiate a line of work focusing on an alternative definition of the RoS constraint, where the constraint is enforced for each auction j

separately rather than the aggregation over all m auctions. Formally, the alternative RoS constraint for each bidding agent i is

$$x_{ij} \cdot v_{ij} \geq \tau_i \cdot p_{ij}, \quad \forall j \in [m]. \quad (\text{RoS}')$$

Under this definition (RoS'), GSP is incentive compatible for the bidding agents, which in general is not the case with (RoS).³

Lv et al. (2023) consider the mechanism design problem with the alternative definition (RoS') when agent with both utility-maximization and value-maximization objectives are present. When their objective types are public, they show that one can use the same efficient allocation rule (higher bids wins higher slots) for all agents and VCG (GSP) payment for utility-maximization (value-maximization) agents. When their objective types are private, they propose a novel mechanism such that the payment of each agent depends on its allocated slot but not their objective type. Under this mechanism, they also prove a 2-approximation in terms of liquid welfare.

6. EMERGING TOPICS

In this section, we cover some emerging topics in the literature that go beyond bidding algorithms, equilibrium and PoA analysis, and optimal auction design.

6.1 Utility functions of advertisers using autobidding

So far this survey has focused on the interaction between bidding agents and the platform (the auctioneer), assuming advertisers' inputs as fixed. However, to fully grasp the impact of auction formats, we must model how advertisers react. In game-theoretic language, most autobidding research has focused on the bidding agent subgame, neglecting the multi-period game where advertisers first submit inputs, followed by the subgame with the bidding agent decisions where the allocation and payment accrues.

The key question in modeling advertiser decisions is whether they are utility-maximizing, value-maximizing or something else. Auction design has traditionally assumed utility maximization, but the rise of target-based bidding strategies challenges this. Why would a utility maximizer use a value-maximizing bidding agent? If instead advertisers' objective is to maximize value subject to a constraint, what incentives guide their input decisions to the autobidding agent?

Regarding the first question, one informal argument for value-maximization agents being favored in practice is a principal-agent model (Fadaei and Bichler, 2017; Bichler and Paulsen, 2018). In this model, each advertiser has a decision department (the principal) and an execution department (the agent) with slightly misaligned goals. To mitigate risk and ensure performance, the principal often sets value-maximization goals for the agent with clear constraints.

Recently, Perloff and Mehta (2023) demonstrate that a utility-maximizing agent prefers to bid through a target-based bidding agent rather than through a marginal-based bidding agent when the platform *lacks commitment* to the declared auction rules: that is, the platform can revisit the rules of the auction (e.g., may readjust

³(RoS) and (RoS') are equivalent when there is only one auction ($m = 1$) and the constraint is ex-post.

reserve prices depending on the bids submitted by the bidders) after bids have already been submitted. Furthermore, they show that due to the lack of commitment the bid shading effect when advertisers bid using a marginal bidding agent is so aggressive that if the platform would enforce to bid only through a marginal bidding agent (e.g. by removing the option of using a target-based autobidder), the platform's revenue would be lower than the revenue they obtain when advertisers use a target-based bidding agent. Bergemann et al. (2023) study the welfare and pricing implications when profit-maximizing advertisers use autobidding systems and lack user data which is known to the autobidder/platform. Compared to the case where advertisers can directly bid in each auction (and all user data is known to them), they show that the autobidding system create negative externalities on external advertising channels (outside of the platform) both in terms of allocation efficiency and consumer surplus.

If in turn advertiser's objectives are aligned with a value-maximizing objective, Alimohammadi et al. (2023) study what type of auctions are *autobidding incentive compatible (AIC)*: for what type of auctions an advertiser with a target-based preference (or a budget-based preference) prefers to submit their constraint as their input to the autobidding agent. They show the second price auction is not AIC for both the target and budget case. For first-price auctions, when bidding agents are restricted to use a uniform policy the auction is AIC, while when they can also use non-uniform bidding strategies then auction is not AIC. More recently, Feng et al. (2023) investigate the PoA of running first-price auctions with budget-constrained autobidders when the budget constraints are strategically chosen by the advertisers and demonstrate constant PoA for such a game.

In addition, there is a second stream of literature on the multi-channel auction problem where they study how value-maximizing advertisers strategically submit their inputs to multiple autobidder agents, where each autobidder agent bids on advertiser's behalf for a particular channel. The following section presents the most interesting results on this topic.

6.2 Multi-channel

In practice, advertisers may procure ad impressions simultaneously on multiple advertising channels. This can involve optimizing campaigns across a single platform's various channels (e.g., Google Ads inventory, including YouTube, Display, Search, Discover, Gmail, and Maps) or across channels owned by different platforms (such as Google, Meta, and Microsoft). In such scenarios, if advertisers are value-maximizing agents subject to a global (ROS) and (BUDGET) then the advertiser's bidding problem and the channel's auction design problem are far from trivial as the advertisers' global constraints interlinks the bidding problem (and, hence, the auction design) across channels.

In what follows, we present recent research that has been trying to shed light on this topic both from an advertiser perspective on how to bid across the channels as well as from a channel perspective on the design of auctions.

6.2.1 Bidding with multiple channels. Deng et al. (2023) study the problem of multi-channel bidding where an advertiser aims to maximize their total conversion while satisfying aggregate (ROS) and (BUDGET) constraints across all channels.

In particular, the advertiser can only utilize two levers on each channel to set up their campaigns, namely setting a per-channel budget and per-channel target RoS. Deng et al. (2023) first analyze the effectiveness of each of these levers via comparison against the global optimum in which the advertiser can directly bid on each impression, and show that: when an advertiser only optimizes over per-channel RoSs, their total conversion can be arbitrarily worse than what they could have obtained in the global optimum, while the advertiser can achieve the global optimum leveraging per-channel budgets only. Under a bandit feedback setting, Deng et al. (2023) further present an efficient and low-regret learning algorithm that produces per-channel budgets whose resulting conversion approximates that of the global optimum. Susan et al. (2023) present a strategy for multi-channel bidding when channels adopt auction rules that may or may not be incentive-compatible under the presence of budget constraints. Aggarwal et al. (2024) characterize the optimal bidding for a continuous query-model where the size of a query is infinitesimal. They show that the advertiser's bidding problem is equivalent to finding a per channel uniform bid such that the advertiser's marginal cost-per-acquisition in each of the channels is the same.

6.2.2 Multi-channel auction design. Aggarwal et al. (2023) initiate the study of multi-channel autobidding auction design focusing on the case of a platform owning multiple internal advertising channels (e.g., Google: Search, Play, YouTube; Meta: Instagram, Facebook, Messenger, etc.) In their setting, they allow a general advertising ecosystem with advertisers having either a (RoS) or (BUDGET) global constraint as well as profit-maximizing advertisers but restrict channels to sell their inventory using a SPA with a reserve price. They study the revenue implications for the platform of having each channel to independently optimize their reserve prices (local optimization) compared to having a global reserve price policy across the channels (global optimization). They consider two models: one in which the channels have full freedom to set reserve prices, and another in which the channels have to respect floor prices set by the publisher. They show that in the first model, welfare and revenue loss from local optimization is bounded by a function of the advertisers' inputs, but is independent of the number of channels and bidders (see Theorem 3 on Aggarwal et al. (2023) for details on the specific bounds). For the second model, they show that the revenue from local optimization could be arbitrarily smaller than those from global optimization.

Aggarwal et al. (2024) study the problem of auction design in the multi-channel setting where multiple platforms (each own a single channel) are competing to sell their inventory to the same pool of advertisers. They consider value-maximizing advertisers that have a (RoS) constraint across channels. The advertisers strategically report target ROIs to each channel's autobidder, which bids uniformly⁴ on their behalf into the channel's auction. Each platform chooses between using a first-price auction or a second-price auction to maximize its own revenue. They show that for a revenue-maximizing platform, competition is a key factor to consider

⁴Note that, while uniform bidding is optimal when each channel is using a truthful auction Aggarwal et al. (2019), uniform bidding is generally not optimal when the channel is running a first-price auction. In this paper, uniform bidding is used to model a practical constraint that a system might impose.

when designing auctions. While first-price auctions are optimal (for both revenue and welfare) in the absence of competition (Deng et al., 2021), this no longer holds true in multi-channel scenarios. Aggarwal et al. (2024) show that for the case of two competing platforms, there exists a large class of valuations for the advertisers such that from the platform’s perspective, running a second-price auction (rather than a first-price auction) is a dominant strategy. They also identify some key factors influencing the platform’s choice of auction format: (i) advertiser sensitivity to price changes – how much the advertisers’ reported targets change against auction changes, (ii) intensity of competition among advertisers, and (iii) relative inefficiency of second-price auctions compared to first-price auctions.

6.3 Empirical Studies

In the previous sections, we discussed the theoretical understanding of autobidding auctions in different aspects. However, the performance of different auction formats is usually analyzed in terms of PoA, which essentially focuses on welfare analysis in the worst-case scenarios, while the real-world instances could have much better equilibrium welfare. To complement the theoretical analysis, Deng et al. (2024) empirically study how different auction formats (namely VCG, FPA and GSP) perform in the autobidding world with synthetic datasets when advertisers adopt different bidding algorithms.

Non-uniform bid scaling. Aggarwal et al. (2019) demonstrate that uniform bid-scaling (i.e., always bid κv with a universal bid-scaling factor κ when the bidder’s value is v) is an optimal strategy for value maximizers in auctions that are truthful for quasi-linear utility maximizers. Therefore, each autobidding agent is only required to optimize one bid-scaling factor to find the best strategy. On the other hand, for auctions that are not truthful for quasi-linear utility maximizers (such as FPA and GSP), uniform bid-scaling can result in a suboptimal bidding strategy, while non-uniform bid-scaling (i.e., use different bid-scaling factors in different auctions) may greatly improve the bidding performance.

Synthetic datasets and experiment setup. To generate the datasets that mimic the data structure of practical ad auctions, Deng et al. (2024) randomly draw query features and bidder features from multidimensional Gaussian distributions, and the bidder’s values are drawn from log-normal distributions parameterized by query and bidder features. To facilitate the simulation of non-uniform bid-scaling algorithms, Deng et al. (2024) partition the queries to different categories following a multi-layer laminar structure. Each bidder chooses different bidding multipliers for different query clusters, and updates the multipliers through a gradient-descent based algorithm in each round.

Empirical Results. When bidders only adopt uniform bid-scaling strategies, it is observed that $\text{FPA} > \text{GSP} > \text{VCG}$ for both welfare and profit. Such a result is consistent with the theoretical finding in the sense that FPA has better welfare and profit (Deng et al., 2021). When bidders can adopt non-uniform bid-scaling strategies, the empirical result of $\text{FPA} > \text{GSP} > \text{VCG}$ for both welfare and profit continues to hold. For different levels of non-uniform bid-scaling algorithms, where a higher non-uniform bid-scaling level corresponds to a larger number of query

clusters with different bid multipliers, there are different trends for different auction formats. For FPA, both profit and welfare decrease as the non-uniform bid-scaling level increases. On the other hand, for GSP, increasing the non-uniform bid-scaling level increases profit but decreases welfare; and for VCG, switching to different levels of non-uniform bid-scaling has no effect on welfare and profit.

7. CONCLUSION

In this survey, we covered a large portion of recent works related to autobidding in the online advertising ecosystem. We mentioned bidding algorithms for both truthful and non-truthful auctions in the presence of RoS and budget constraints. We discussed the existence of equilibrium, the price of anarchy with respect to different solution concepts, and the convergence properties of several bidding dynamics. We introduced recent advancements in terms of revenue-optimal auction design under different information structures and with various benchmarks. Finally, we discussed emerging topics in the literature, such as the role of advertiser decision, the application with multi-channel, and the comparison between theoretical and empirical results. We hope this survey provides a valuable resource for both practitioners and academics seeking to understand the state-of-the-art in this rapidly evolving field.

REFERENCES

- AGGARWAL, G., BADANIDIYURU, A., AND MEHTA, A. 2019. Autobidding with constraints. In *Web and Internet Economics: 15th International Conference, WINE 2019, New York, NY, USA, December 10–12, 2019, Proceedings 15*. Springer, 17–30.
- AGGARWAL, G., FIKIORIS, G., AND ZHAO, M. 2024. No-regret algorithms in non-truthful auctions with budget and roi constraints. *arXiv preprint arXiv:2404.09832*.
- AGGARWAL, G., PERLROTH, A., SCHVARTZMAN, A., AND ZHAO, M. 2024. Platform competition in the autobidding world. *arXiv preprint arXiv:2405.02699*.
- AGGARWAL, G., PERLROTH, A., AND ZHAO, J. 2023. Multi-channel auction design in the autobidding world. In *Proceedings of the 24th ACM Conference on Economics and Computation*. EC '23. Association for Computing Machinery, New York, NY, USA, 21.
- AI, R., WANG, C., LI, C., ZHANG, J., HUANG, W., AND DENG, X. 2022. No-regret learning in repeated first-price auctions with budget constraints. *arXiv preprint arXiv:2205.14572*.
- ALIMOHAMMADI, Y., MEHTA, A., AND PERLROTH, A. 2023. Incentive compatibility in the auto-bidding world. In *Proceedings of the 24th ACM Conference on Economics and Computation*. EC '23. Association for Computing Machinery, New York, NY, USA, 63.
- AUSUBEL, L. M. 2004. An efficient ascending-bid auction for multiple objects. *American Economic Review* 94, 5, 1452–1475.
- BADANIDIYURU, A., LANGFORD, J., AND SLIVKINS, A. 2014. Resourceful contextual bandits. In *Conference on Learning Theory*. PMLR, 1109–1134.

- BALSEIRO, S., DENG, Y., MAO, J., MIRROKNI, V., AND ZUO, S. 2021a. Robust auction design in the auto-bidding world. *Advances in Neural Information Processing Systems* 34, 17777–17788.
- BALSEIRO, S., DENG, Y., MAO, J., MIRROKNI, V., AND ZUO, S. 2024. Optimal mechanisms for a value maximizer: The futility of screening targets. In *Proceedings of the 25th ACM Conference on Economics and Computation*.
- BALSEIRO, S., KIM, A., MAHDIAN, M., AND MIRROKNI, V. 2021. Budget-management strategies in repeated auctions. *Operations Research* 69, 3, 859–876.
- BALSEIRO, S., KROER, C., AND KUMAR, R. 2023. Contextual standard auctions with budgets: Revenue equivalence and efficiency guarantees. *Management Science* 69, 11, 6837–6854.
- BALSEIRO, S. R., BESBES, O., AND WEINTRAUB, G. Y. 2015. Repeated auctions with budgets in ad exchanges: Approximations and design. *Management Science* 61, 4, 864–884.
- BALSEIRO, S. R., BHAWALKAR, K., FENG, Z., LU, H., MIRROKNI, V., SIVAN, B., AND WANG, D. 2023. Joint feedback loop for spend and return-on-spend constraints. *arXiv preprint arXiv:2302.08530*.
- BALSEIRO, S. R., DENG, Y., MAO, J., MIRROKNI, V., AND ZUO, S. 2022. Optimal mechanisms for value maximizers with budget constraints via target clipping. In *Proceedings of the 23rd ACM Conference on Economics and Computation*. 475–475.
- BALSEIRO, S. R., DENG, Y., MAO, J., MIRROKNI, V. S., AND ZUO, S. 2021b. The landscape of auto-bidding auctions: Value versus utility maximization. In *Proceedings of the 22nd ACM Conference on Economics and Computation*. 132–133.
- BALSEIRO, S. R. AND GUR, Y. 2019. Learning in repeated auctions with budgets: Regret minimization and equilibrium. *Management Science* 65, 9, 3952–3968.
- BALSEIRO, S. R., LU, H., AND MIRROKNI, V. 2023. The best of many worlds: Dual mirror descent for online allocation problems. *Operations Research* 71, 1, 101–119.
- BERGEMANN, D., BONATTI, A., AND WU, N. 2023. Managed campaigns and data-augmented auctions for digital advertising. In *Proceedings of the 24th ACM Conference on Economics and Computation*. EC '23. Association for Computing Machinery, New York, NY, USA, 271.
- BICHLER, M. AND PAULSEN, P. 2018. A principal-agent model of bidding firms in multi-unit auctions. *Games and Economic Behavior* 111, 20–40.
- CASTIGLIONI, M., CELLI, A., AND KROER, C. 2022. Online learning with knapsacks: the best of both worlds. In *International Conference on Machine Learning*. PMLR, 2767–2783.
- CASTIGLIONI, M., CELLI, A., AND KROER, C. 2023. Online learning under budget and roi constraints via weak adaptivity. *arXiv preprint arXiv:2302.01203*.
- CHEN, X., KROER, C., AND KUMAR, R. 2021. The complexity of pacing for second-price auctions. In *EC (2021-07-22)*.

- CONITZER, V., KROER, C., SODOMKA, E., AND STIER-MOSES, N. E. 2022. Multiplicative pacing equilibria in auction markets. *Oper. Res.* 70, 2 (mar), 963–989.
- DENG, Y., GOLREZAEI, N., JAILLET, P., LIANG, J. C. N., AND MIRROKNI, V. 2022. Individual welfare guarantees in the autobidding world with machine-learned advice. *arXiv preprint arXiv:2209.04748*.
- DENG, Y., GOLREZAEI, N., JAILLET, P., LIANG, J. C. N., AND MIRROKNI, V. 2023. Multi-channel autobidding with budget and roi constraints. In *International Conference on Machine Learning*. PMLR.
- DENG, Y., MAHDIAN, M., MAO, J., MIRROKNI, V., ZHANG, H., AND ZUO, S. 2024. Efficiency of the generalized second-price auction for value maximizers. In *Proceedings of the ACM on Web Conference 2024*. WWW’24. Association for Computing Machinery, 46–56.
- DENG, Y., MAO, J., MIRROKNI, V., TENG, Y., AND ZUO, S. 2024. Non-uniform bid-scaling and equilibria for different auctions: An empirical study. In *Proceedings of the ACM on Web Conference 2024*. WWW’24. Association for Computing Machinery, 256–266.
- DENG, Y., MAO, J., MIRROKNI, V., ZHANG, H., AND ZUO, S. 2022. Efficiency of the first-price auction in the autobidding world. *arXiv preprint arXiv:2208.10650*.
- DENG, Y., MAO, J., MIRROKNI, V., ZHANG, H., AND ZUO, S. 2023. Autobidding auctions in the presence of user costs. In *Proceedings of the ACM Web Conference 2023*. 3428–3435.
- DENG, Y., MAO, J., MIRROKNI, V., AND ZUO, S. 2021. Towards efficient auctions in an auto-bidding world. In *Proceedings of the Web Conference 2021*. 3965–3973.
- DOBZINSKI, S. AND LEME, R. P. 2014. Efficiency guarantees in auctions with budgets. In *International Colloquium on Automata, Languages, and Programming*. Springer, 392–404.
- FADAEI, S. AND BICHLER, M. 2017. Truthfulness with value-maximizing bidders: On the limits of approximation in combinatorial markets. *European Journal of Operational Research* 260, 2, 767–777.
- FENG, Y., LUCIER, B., AND SLIVKINS, A. 2023. Strategic budget selection in a competitive autobidding world. *arXiv preprint arXiv:2307.07374*.
- FENG, Z., PADMANABHAN, S., AND WANG, D. 2023. Online bidding algorithms for return-on-spend constrained advertisers. In *Proceedings of the ACM Web Conference 2023*. 3550–3560.
- GOEL, G., MIRROKNI, V., AND PAES LEME, R. 2014. Clinching auctions beyond hard budget constraints. In *Proceedings of the fifteenth ACM conference on Economics and computation*. 167–184.
- GOLREZAEI, N., LOBEL, I., AND PAES LEME, R. 2021. Auction design for roi-constrained buyers. In *Proceedings of the Web Conference 2021*. 3941–3952.
- HAN, Y., ZHOU, Z., AND WEISSMAN, T. 2024. Optimal no-regret learning in repeated first-price auctions. *Operations Research*.
- HE, X., PAN, J., JIN, O., XU, T., LIU, B., XU, T., SHI, Y., ATALLAH, A., HERBRICH, R., BOWERS, S., ET AL. 2014. Practical lessons from predicting

- clicks on ads at facebook. In *Proceedings of the Eighth International Workshop on Data Mining for Online Advertising*. 1–9.
- JUAN, Y., ZHUANG, Y., CHIN, W.-S., AND LIN, C.-J. 2016. Field-aware factorization machines for ctr prediction. In *Proceedings of the 10th ACM conference on recommender systems*. 43–50.
- LI, J. AND TANG, P. 2024. Vulnerabilities of single-round incentive compatibility in auto-bidding: Theory and evidence from roi-constrained online advertising markets. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI-24*. International Joint Conferences on Artificial Intelligence Organization. Main Track.
- LIANG, J. C. N., LU, H., AND ZHOU, B. 2023. Online ad procurement in non-stationary autobidding worlds. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- LIAW, C., MEHTA, A., AND PERLROTH, A. 2023. Efficiency of non-truthful auctions in auto-bidding: The power of randomization. In *Proceedings of the ACM Web Conference 2023*. WWW '23. Association for Computing Machinery, New York, NY, USA, 3561–3571.
- LIAW, C., MEHTA, A., AND ZHU, W. 2024. Efficiency of non-truthful auctions in auto-bidding with budget constraints. In *Proceedings of the ACM Web Conference 2024*. WWW '24.
- LIU, X. AND SHEN, W. 2023. Auto-bidding with budget and roi constrained buyers. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI 2023, 19th-25th August 2023, Macao, SAR, China*. International Joint Conferences on Artificial Intelligence Organization.
- LU, P., XU, C., AND ZHANG, R. 2023. Auction design for value maximizers with budget and return-on-spend constraints. In *Web and Internet Economics: 19th International Conference, WINE 2023, Shanghai, China, December 4–8, 2023, Proceedings 19*. Springer.
- LU, Q., PAN, S., WANG, L., PAN, J., WAN, F., AND YANG, H. 2017. A practical framework of conversion rate prediction for online display advertising. In *Proceedings of the ADKDD'17*. 1–9.
- LUCIER, B., PATTATHIL, S., SLIVKINS, A., AND ZHANG, M. 2023. Autobidders with budget and roi constraints: Efficiency, regret, and pacing dynamics. *arXiv preprint arXiv:2301.13306*.
- LV, H., BEI, X., ZHENG, Z., AND WU, F. 2023. Auction design for bidders with ex post roi constraints. In *Web and Internet Economics: 19th International Conference, WINE 2023, Shanghai, China, December 4–8, 2023, Proceedings 19*. Springer.
- LV, H., ZHANG, Z., ZHENG, Z., LIU, J., YU, C., LIU, L., CUI, L., AND WU, F. 2023. Utility maximizer or value maximizer: mechanism design for mixed bidders in online advertising. In *Proceedings of the Thirty-Seventh AAAI Conference on Artificial Intelligence and Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence and Thirteenth Symposium on Educational Advances in Artificial Intelligence*. AAAI'23/IAAI'23/EAAI'23. AAAI Press.

- McMAHAN, H. B., HOLT, G., SCULLEY, D., YOUNG, M., EBNER, D., GRADY, J., NIE, L., PHILLIPS, T., DAVYDOV, E., GOLOVIN, D., ET AL. 2013. Ad click prediction: a view from the trenches. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*. 1222–1230.
- MEHTA, A. 2022. Auction design in an auto-bidding setting: Randomization improves efficiency beyond vcg. In *Proceedings of the ACM Web Conference 2022*. 173–181.
- MYERSON, R. B. 1981. Optimal auction design. *Mathematics of operations research* 6, 1, 58–73.
- PAES LEME, R., PILIOURAS, G., SCHNEIDER, J., SPENDLOVE, K., AND ZUO, S. 2024. Complex dynamics in autobidding systems. In *Proceedings of the 25th ACM Conference on Economics and Computation*.
- PERLROTH, A. AND MEHTA, A. 2023. Auctions without commitment in the auto-bidding world. In *Proceedings of the ACM Web Conference 2023*. WWW '23. Association for Computing Machinery, New York, NY, USA, 3478–3488.
- SUSAN, F., GOLREZAEI, N., AND SCHRIJVERS, O. 2023. Multi-platform budget management in ad markets with non-ic auctions. *arXiv preprint arXiv:2306.07352*.
- WANG, Q., YANG, Z., DENG, X., AND KONG, Y. 2023. Learning to bid in repeated first-price auctions with budgets. In *International Conference on Machine Learning*. PMLR, 36494–36513.
- WILKENS, C. A., CAVALLO, R., AND NIAZADEH, R. 2017. Gsp: the cinderella of mechanism design. In *Proceedings of the 26th International Conference on World Wide Web*. 25–32.
- WILKENS, C. A., CAVALLO, R., NIAZADEH, R., AND TAGGART, S. 2016. Mechanism design for value maximizers. *arXiv preprint arXiv:1607.04362*.
- XING, Y., ZHANG, Z., ZHENG, Z., YU, C., XU, J., WU, F., AND CHEN, G. 2023. Truthful auctions for automated bidding in online advertising. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI-23*. International Joint Conferences on Artificial Intelligence Organization, 2915–2922. Main Track.
- ZHOU, G., ZHU, X., SONG, C., FAN, Y., ZHU, H., MA, X., YAN, Y., JIN, J., LI, H., AND GAI, K. 2018. Deep interest network for click-through rate prediction. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1059–1068.
- ZHOU, Y., CHAKRABARTY, D., AND LUKOSE, R. 2008. Budget constrained bidding in keyword auctions and online knapsack problems. In *Proceedings of the 17th international conference on world wide web*. 1243–1244.