Overview. My research is in algorithms, primarily in problems concerned with *incentives*. The ubiquity of the internet has given rise to more and more settings that involve self-interested individuals—be they ride-sharing users, advertisers bidding in an advertising auction, or reporters maximizing the spread of news in a social network. While the majority of my work has contributed to **core problems within algorithms and mechanism design**, my recent and forthcoming work focuses on the topic of **mechanism design for social good**.

For a concrete example of an algorithm that interacts with agents who have a strategic interest in the outcome, take the E.U. Emissions Trading Scheme (ETS). The ETS elicits the economic value produced by each manufacturer and calculates the social cost of carbon pollution. Then, with the objective of trading off value to the economy and the cost of pollution to society, the ETS algorithmically determines which European manufacturers can produce (and how many metric tons of) carbon for the next year. If a manufacturer can misrepresent the economic value of their products such that the government increases their production allowance, leading the manufacturer to earn more revenue, then they will do so—and this behavior can render our classical algorithmic solutions insufficient. Reasoning about participants' incentives complicates the design task, resulting in a new frontier in optimization and algorithm design. The field of *algorithmic mechanism design* uses tools from game theory and economics to design *mechanisms*: algorithms that guarantee that *even when the participants act in their own self-interest*, the objective (i.e., a good allocation of carbon emission licenses) is achieved.

Within the foundations of mechanism design, my research is centered in *multi-parameter* settings, where participants have multiple pieces of private information impacting their interests, and algorithmic solutions sit at the threshold of tractability. A manufacturer's economic value is multi-parameter—for each metric ton of carbon emitted, more goods are produced, and supply eventually exceeds demand, decreasing value—thus the firm has a different value for each additional metric ton emitted. **My work has given the first positive results for long-standing open problems:**

- In settings where varying quality levels of a service are sold, we provide the *first explicit characterization* of revenue-optimal auctions for any general multi-parameter setting [12, 18].
- In the interdependent values model, where a buyer's valuation of an item depends not only on their own private information, but on other buyers' information as well, we provide the *first* approximation to optimal *welfare*—the sum of players' values for the items they are awarded—for multi-dimensional and even combinatorial settings [15, 16]. This work was awarded the **EC 2019 Best Paper with a Student Lead Author**.
- In a market where buyers and sellers trade, we provide the *first* approximation to the gains from trade—the value of the efficient trades—for a setting with a *multi-parameter buyer* [7].

By attacking problems from the most complex domains, my work not only provides algorithmic solutions, but develops cutting-edge mathematical theory.

While designing mechanisms, **I use *simplicity* and *robustness* as guiding principles**. *Simple* mechanisms are easy to implement and easy for participants to understand; complicated mechanisms may be inaccurately predicted by game theory due to participants' inability to understand how to strategize [3]. By designing a mechanism to be *robust*—to lack of detailed information, agents' risk attitudes, or agents' valuation classes—we expand the settings for which its guarantees hold. Otherwise, a small amount of noise in a dataset or model could destroy theoretical guarantees. By possessing these features, my work provides tangible guidelines for use in practice.

While many of my results have focused on classical auction design, my recent work applies insights and mathematical tools from algorithmic mechanism design to develop the foundations of *mechanism design for social good*, particularly in domains such as healthcare, environmental sustainability, and online labor markets. Each of these domains is concerned with how the government and citizens design allocation policies, impose tax structures, create laws, and regulate activities; *these are mechanism design problems.* These topics have been well-explored in economics, operations research, sociology, public policy, and many other adjacent fields, yet have been under-explored in computer science. Because of this, with Rediet Abebe in 2016, **I co-founded the Mechanism Design for Social Good initiative** [1], a multi-institutional interdisciplinary research initiative focused on improving access to opportunity for historically underserved and disadvantaged communities. What has resulted is an interdisciplinary community of over 2300 researchers globally, toward which I contribute both leadership (e.g., organizing, PC-chairing, and steering the four workshops and first ACM MD4SG conference) as well as research [2, 4, 7, 17, 19] (e.g., determining parameters that guarantee the carbon ETS used in practice is robust to strategic behavior [19]).

Road Map. The next section describes my work in mechanism design for social good. The following section details my work within foundational mechanism design, e.g., uncovering a *spectrum* of auction complexity, and using simple mechanisms to guarantee robustness to misinformation or unexpected behavior. Finally, I discuss my long-term research agenda.

Mechanism Design for Social Good

Theoretical mechanism design has primarily guided both the foundations and implementations of auctions; most of my thesis research (described in the next section) has focused on core directions within this domain. My goal is to extend this rich mathematical theory toward *social good* domains. Of course, there are many different aspects to social good. I am particularly interested in algorithmic questions regarding mitigation of social harms where incentives must be considered.

Carbon Pricing. In energy markets, firms produce and sell goods, but in the production of said goods, they emit pollution, a consequence with a negative impact on society. A government designs their Emissions Trading Scheme (ETS) with a goal to maximize welfare: the total economic value from the firms selling goods minus the societal cost of pollution. In work with co-authors [19], we investigate the mechanism most used in Emissions Trading Schemes globally: the uniform price auction [11] with a *cap* on how many metric tons can be emitted and a *minimum price* that polluters must pay per license. This auction is not truthful. It sets the license price at the maximum of the lowest-winning-bid and the minimum price, so firms can under-bid to potentially cause a lower price to be set. If successful, the result is that firms producing less economic value pollute in place of under-bidding firms who could have produced more economic value, hence more carbon is emitted than is made up for in economic gain, resulting in an unbounded *price of anarchy.*[1] We determine that by setting a "safe" minimum price to be a particular function of the cap, we can recover strong price of anarchy guarantees. This means that using a safe price ensures approximately the best welfare of *this* uniform-price auction (for this cap and safe price) from non-strategic agents. However, it does not guarantee an approximation to *optimal welfare* over *any* uniform-price auction with any choice of cap and price. We prove that using a safe price with the uniform-price auction either achieves near-optimal welfare, or, there must be some single dominant manufacturer, and the government should instead contract solely with them to achieve near-optimal welfare.

Health Insurance. In the US, there are two common types of health insurance markets: (1) free markets with unregulated entry, e.g., the Affordable Care Act Exchanges, where any insurance provider can join and offer a plan at any price, and (2) regulated markets, e.g., within a corporation, that limit entry into the market via some initial mechanism, and the employees of the corporation then purchase a plan from among the chosen insurance providers. Compared to unregulated markets with full choice over plans, regulated markets limit entry to the k plans with the cheapest prices, inherently removing some of the choices (decreasing consumer utility) yet also forcing lower prices (increasing utility). In our work [17] we ask: is limiting entry better for consumer *utility*? To answer this question, we must first characterize insurance provider prices in equilibrium—a typically computationally intensive and mechanical mathematical process. However, we observe that, surprisingly, we can leverage techniques from revenue maximization to solve for these prices in an *intuitive* way. Then, we compare the utility of both settings, and again leverage a theorem from

[1] The price of anarchy is the ratio of the optimal (non-strategic) allocation to the worst-case equilibrium.

the foundations of revenue maximization to transform the condition into a precise characterization of when limiting entry improves consumer utility. This work not only gives insight into health insurance market regulation, but also serves as a proof of concept as to how foundational mechanism design can find applications in social good!

Online Labor Markets. In online labor markets such as Upwork or Mechanical Turk, workers form one side of the platform, selling their labor, and employers form the other side, purchasing labor. The platforms employ algorithms to aid in the matchmaking process, and their algorithms have tangible impacts on employment opportunities, the quality of work received, and the incentives of both parties to report information truthfully. To ensure that a platform is benefitting its users, we can ask what *gains from trade*, or quality of matches, its algorithm achieves. We know from a seminal result [29] that it is provably impossible to achieve the optimal gains from trade while satisfying natural incentive and no-deficit constraints. In one approach [4], my co-authors and I take a resource augmentation perspective: recruiting additional employers to the platform and proving that the gains from trade of a very simple mechanism on this larger market (with increased competition) beat the impossible-to-achieve optimal gains from trade of the original market. The mechanism we use, adapted from McAfee [26], is robust, as it does not rely on market information.

If instead our market size is fixed—we cannot recruit additional employers or workers—then we can instead *approximate* the optimal gains from trade. In very recent joint work [7], we design simple mechanisms that give the first gains from trade approximation guarantee for a *multi-parameter* employer, that is, one who is interested in hiring multiple types of workers.

Foundations of Mechanism Design

Interdimensional Settings: An Auction Complexity Middle Ground

The problem of designing revenue-maximizing auctions has been the subject of intense study for decades. The problem is completely solved when the auctioneer has a single item to sell [28]; however, even when the seller has only two items to sell, optimal auctions are not well understood. The optimal auction in this setting can be incredibly complex—highly randomized and intractable to compute [10, 25]—and thus characterizations remain elusive.

In our work [12, 18], we uncover a whole *spectrum* of complexity between the the "easy" case of "single-parameter" settings (where the buyer's private information is characterized by a single number) and the fully general "multi-dimensional" setting. We introduce a new subclass of natural and ubiquitous multi-dimensional settings that we call *interdimensional* settings, for which the complexity of maximizing revenue occupies this complex-yet-tractable middle ground.

Our main result is a *complete* characterization of the optimal solution in the interdimensional setting we call "The Fedex Problem," where a buyer has a value and a deadline, and a shipping service is worth their value to them only if their package is received on or before their deadline [18]. This is perhaps the simplest multi-dimensional setting imaginable, and the *total ordering* over shipping options—that 1-day shipping is better than 2-day shipping and so on—greatly simplifies the buyer's incentive constraints. As a result, a dual formulation of the problem has fewer variables, making it more tractable to solve. We obtain an explicit closed-form solution for the dual variables, and then algorithmically reconstruct the primal variables (the optimal mechanism); this is the first optimal characterization for a multi-dimensional setting without any restriction on the form of the mechanism or on the buyer value distributions. En route, we also develop multi-dimensional analogues of Myerson's Nobel-Prize-winning theory: revenue curves, ironing (needed even in the single-bidder case), and virtual values.

We then characterize the optimal mechanism for selling to a single-minded buyer, for whom a set of items is worth their value only if it contains their desired bundle—bundles are now *partially-ordered* by containment. The single-minded setting was one of the first settings studied within

Table 1: The complexity spectrum from interdimensional settings; bold results are my own [12, 18].

Setting / Metric	1 item	FedEx	Budgets	Single-Minded	Multi-Unit	Coordinated Valuations	2 items
Characterization	closed form [28]	**explicit dual**	dual [14] properties	**dual properties**	—	—	intractable [10, 25]
Menu Size LB	1 [28]	exponential [30]	exponential [14]	**unbounded**[2]	**unbounded**	**countable**	uncountable [10, 25]
Menu Size UB		**exponential**		**finite**[3]	—		

algorithmic game theory, demonstrating the limits of polynomial time computation for welfare approximation [24]. We discover that the optimal auction's *menu complexity*, or the number of distinct options offered to the buyer, is *unbounded*[2] *but finite*[3] [12]. This sharply separates the single-minded setting from both the FedEx setting (exponential menu complexity [18, 30]) and the "multi"-dimensional two-item additive setting (uncountable menu complexity [10, 25]).

My work demonstrates that what was previously believed to be a dichotomy of single- and "multi"-dimensional settings is in fact a full complexity spectrum—both by menu size and the optimal mechanism's characterization. Through my work and the follow-up work in additional interdimensional domains [13, 14, 30, 31], we now understand the spectrum in Table 1.

The Efficacy of Simple and Robust Mechanisms

In order to achieve optimal revenue, optimal mechanisms are *precisely tuned to their input*, often resulting in complex mechanisms that, if the input varies even slightly, will not sustain their guarantees. For example, an *uncertainty-averse* buyer is one who discounts uncertain outcomes, and thus pays a premium for outcomes that are more certain. If such a buyer were to participate in the highly randomized optimal mechanism described for the FedEx setting, they might purchase 1-day shipping (which is never randomized) if they can afford it, and otherwise not purchase any shipping option—even though their deadline may be much later than 1 day from now. As a result, the optimal *risk-neutral* mechanism for FedEx fails to earn revenue from many *uncertainty-averse* customers who are unwilling to purchase randomized options, causing it to earn quite poor revenue. In contrast, simple mechanisms—such as posted prices or selling the grand bundle of all items—are more robust to a change in buyer risk attitude or value distribution. In my work, various co-authors and I design mechanisms of these sorts and prove that they guarantee good approximations to many objectives (welfare [8, 15, 16], revenue [6, 8, 9, 15, 16, 20], and gains from trade [4, 7]) independent of (1) seller knowledge of detailed buyer information [4, 15, 16, 20], (2) buyer risk attitudes [9], (3) buyer valuation classes [6] (e.g., with or without complementarities).

Interdependent Buyers. In the standard model of *independent private values*, a buyer's value for an item is unaffected by other buyers' values for the item. This is often unrealistic. For an item such as a house or a painting, the item's resale value is directly relevant to the buyer's willingness-to-pay, and is well-approximated by the values of other buyers. In many cases, buyers do not actually know their values for an item, but rather, each have partial information regarding the item. Combining the information of all buyers informs a buyer of their value for the item. For example, if two firms are bidding for oil drilling rights, and one firm learns the amount of oil available, that information directly impacts the value that the other firm has for the drilling rights. The *interdependent values model* [27] captures the idea that a buyer's information and value may impact the values of other buyers.

What makes these settings so challenging (and worthy of the 2020 Nobel Prize in Economics) is that changing one's bid impacts the values of *other* bidders. For example, a kindergarten dad

[2]For every M, there exists a prior distribution for which the menu complexity of the optimal mechanism is $\geq M$.
[3]For every prior distribution, the menu complexity is finite.

knows whether a drawing of a dog is by Picasso or his daughter. Upon reporting the low signal (daughter), he wins the drawing, as he is the only one who values the doodle; for reporting the high signal (Picasso), his bid increases the value of art collectors, whose values dwarf his own. Thus, by *increasing* his bid, he *decreases* his chances of winning (as this increases another bidder's value).

To avoid this phenomenon, *all* prior work assumes that valuations satisfy the restrictive "single-crossing" condition—a condition ensuring that an increased bid can only increase chances of winning. However, (1) there is no generalization of this condition past the single-parameter setting, and (2) even simple examples like the dog drawing above are not covered by this condition, and for such examples, *no prior work gives any guidance* on how to allocate! For giving the first (constant-factor) approximation without single-crossing [15, 16] and the *first* welfare guarantee for *any* non-degenerate multi-parameter setting, including up to general combinatorial auctions [16], our work was awarded **EC 2019 Best Paper with a Student Lead Author**.

Future Work

I am eager to continue work on algorithmic problems inspired by society and policy where incentive issues arise. As a theoretical computer scientist, I find it crucial to engage in interdisciplinary conversations in order to understand how and where domain experts believe my research methodology can be used to provide insight on societal problems. This is what provides the essential component of ensuring that my work is of interest to those in the domain, rather than isolated in theory. I am looking to collaborate with experts in and outside of computer science—e.g., with those in security and privacy, cryptography, machine learning, and computing for development, as well as with medical centers, pedagogical scholars, and legal scholars. The following are a few examples of a long-term agenda in mechanism design for social good.

Healthcare. Imagine a number of Medicaid patients who need to receive a medical procedure. The government covers the cost, but it cannot afford to give each of the patients their preferred hospital, and thus must choose how to assign patients to hospitals. Typically, to elicit true preferences, we use payments, but this is infeasible when we are treating impoverished patients. Instead, payments are implemented as wait times or reductions in quality of care; we thus of course prefer to minimally use such payments. The mechanism designer's goal is then to maximize the sum of the patients' true utilities: allocated value minus "payment" (i.e., wait time), subject to the constraints that the agents are incentivized to truthfully report their needs, and that the budget constraints of the government are met. This "money burning" setting where the "payment" does not benefit anyone and is thus "burnt," while resolved for allocating a single treatment [5, 21], is wide open in the multi-dimensional setting of allocating differing hospitals. Within foundational mechanism design, we have a rich toolkit for maximizing patient value (welfare) or focusing on payments (revenue), yet our techniques do not apply to the challenging task of simultaneously maximizing patient value *and* minimizing payments. En route to finding a solution that would give insight for allocating treatment, we must develop fresh ideas—techniques for obtaining a benchmark, potential simple mechanisms, candidate duals—to handle this unusual objective.

Mechanisms and Law. Equal Employment Opportunity (EEO) laws require that protected classifications (e.g., race, national origin, age, disability, gender, sexuality, and religion) cannot be used in hiring, firing, or the terms and conditions of one's employment. One application where labor laws are relevant is contract design. From an economic standpoint, this is the design of payments from employers to workers based *not* on an agent's costly labor, but only on the observable *outcome* of some unknown actions of labor. From an algorithmic standpoint, contract design is the setting where the algorithm designer knows the input to the algorithm and the objective, but the algorithm itself is performed by strategic agents with *no* stake in the outcome, for whom taking actions on behalf of the algorithm is costly. The designer then uses payments *to* the agents in order to incentivize them to take actions that will produce the desired outcome. By EEO laws, payment for equivalent work cannot be different for workers of varying protected classes—even though different workers might incur different costs. How can we design approximately-optimal contracts that adhere to EEO labor laws?

Privacy and Terms of Service Agreements. In order to use most internet applications, a user is presented with a non-negotiable Terms of Service agreement that they must opt into, including agreeing to share their data. Regardless of the application, such agreements are almost always bilateral between only the user and the data-collector. However, one user sharing their data impacts more than just themself—it may impact how someone else is classified or what that person is charged. In ongoing work with Tim Roughgarden, privacy scholar Helen Nissenbaum, and legal scholar Salomé Viljoen, we are studying the impact of using bilateral agreements on incentives and social welfare compared to alternative privacy agreements.

Privacy and Data-Dependent Markets. The E.U.'s General Data Protection Regulation (GDPR) dictates consumers' privacy rights (to opt-in to, erase, and transfer their data) and firms' requirements for data security. Very recent work [23] studies a simple market with a personalizable product, such as a generic radio station vs. a trained Pandora station, and under what conditions consumer utility and/or firm profit are improved when the market is subject to GDPR regulations. However, their work assumes that upon opting into data sharing, a user's true data is simply known, as opposed to elicited. Given that a product is personalized *over time*, such as a pandora station being trained by likes and dislikes, one should expect a consumer to potentially misreport their data—whether directly or, e.g., via using a VPN to provide an alternative location—*especially* if their data impacts the price at which the good is offered to them. Multi-stage market design subject to privacy regulations is a wide-open and important direction that I am eager to work on.

Diverse Admissions. One issue that we face as academics when we run graduate admissions in the US is a gap in our ability to understand the applications coming from varying backgrounds. When an applicant comes from a top research university and has a letter from a well-known colleague, then the signal from their application is very close to our actual value for the student. In contrast, when a student comes from a small liberal arts college, or a historically black college, or a university from an unfamiliar country, we may not know how to evaluate their coursework and research experience as accurately, so the application may provide a much noisier signal of our value. Is it possible that it's not only better for society and breadth of perspectives, but also for pure technical quality to invest in costly inspection procedures in order to better estimate and thus accept the values of applicants from diverse backgrounds? These questions, which have arisen in my own admissions experiences, may be well-modeled by the information acquisition literature, such as a matching setting with costly interviews [22] or a Pandora's box problem [32]. In addition, because faculty are responsible for the students that they admit—with respect to time, finances, and career trajectory—then they bear some *risk*, so it makes sense to investigate risk attitudes.

Conclusion

My research has solved longstanding open problems in many difficult multi-parameter settings, including characterizations of revenue-optimal mechanisms, welfare guarantees in the interdependent values model, approximations to gains from trade. My research has also pioneered the direction of "mechanism design for social good," a community which I co-founded and has since massively grown. I am eager to continue pursuit of the above research program, applying my expertise from multi-parameter algorithmic mechanism design to societal problems where these techniques are likely to provide guidance. I look forward to interfacing with experts in other disciplines in order to ensure that the problems I solve are not only technically interesting, but also relevant and practical. I also intend to select problems that I know will be of special interest to the mechanism design community, bringing them along on this new agenda.

References

[1] Rediet Abebe and Kira Goldner. The Mechanism Design for Social Good Initiative. `http://md4sg.com`, 2016.

[2] Rediet Abebe and Kira Goldner. Mechanism design for social good. *AI Matters*, 4(3):27–34, 2018.

[3] Maurice Allais. Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école américaine. *Econometrica: Journal of the Econometric Society*, pages 503–546, 1953.

[4] Moshe Babaioff, Kira Goldner, and Yannai A. Gonczarowski. Bulow-klemperer-style results for welfare maximization in two-sided markets. In Shuchi Chawla, editor, *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*, pages 2452–2471. SIAM, 2020.

[5] Mark Braverman, Jing Chen, and Sampath Kannan. Optimal provision-after-wait in healthcare. *Mathematics of Operations Research*, 41(1):352–376, 2016.

[6] Yang Cai, Nikhil R. Devanur, Kira Goldner, and R. Preston McAfee. Simple and approximately optimal pricing for proportional complementarities. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, EC '19, Phoenix, AZ, USA, 2019. ACM.

[7] Yang Cai, Kira Goldner, Steven Ma, and Mingfei Zhao. On Multi-Dimensional Gains from Trade Maximization, 2020.

[8] Shuchi Chawla, Kira Goldner, Anna R. Karlin, and J. Benjamin Miller. Non-adaptive matroid prophet inequalities. In *Manuscript*, 2020.

[9] Shuchi Chawla, Kira Goldner, J. Benjamin Miller, and Emmanouil Pountourakis. Revenue maximization with an uncertainty-averse buyer. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '18, pages 2050–2068, USA, 2018. Society for Industrial and Applied Mathematics.

[10] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. Strong duality for a multiple-good monopolist. *Econometrica*, 85(3):735–767, 2017.

[11] Bart de Keijzer, Evangelos Markakis, Guido Schäfer, and Orestis Telelis. Inefficiency of standard multi-unit auctions. In Hans L. Bodlaender and Giuseppe F. Italiano, editors, *Algorithms – ESA 2013*, pages 385–396, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.

[12] Nikhil R. Devanur, Kira Goldner, Raghuvansh R. Saxena, Ariel Schvartzman, and S. Matthew Weinberg. Optimal mechanism design for single-minded agents. In Péter Biró, Jason Hartline, Michael Ostrovsky, and Ariel D. Procaccia, editors, *EC '20: The 21st ACM Conference on Economics and Computation, Virtual Event, Hungary, July 13-17, 2020*, pages 193–256. ACM, 2020.

[13] Nikhil R. Devanur, Nima Haghpanah, and Christos-Alexandros Psomas. Optimal Multi-Unit Mechanisms with Private Demands. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, EC '17, pages 41–42, New York, NY, USA, 2017. ACM.

[14] Nikhil R. Devanur and S. Matthew Weinberg. The Optimal Mechanism for Selling to a Budget Constrained Buyer: The General Case. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, EC '17, pages 39–40, New York, NY, USA, 2017. ACM.

[15] Alon Eden, Michal Feldman, Amos Fiat, and Kira Goldner. Interdependent values without single-crossing. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, EC '18, pages 369–369, New York, NY, USA, 2018. ACM.

[16] Alon Eden, Michal Feldman, Amos Fiat, Kira Goldner, and Anna R. Karlin. Combinatorial auctions with interdependent valuations: Sos to the rescue. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, EC '19, Phoenix, AZ, USA, 2019. ACM.

[17] Meryem Essaidi, Kira Goldner, and S. Matthew Weinberg. When to limit market entry under mandatory purchase. *CoRR*, abs/2002.06326, 2020.

[18] Amos Fiat, Kira Goldner, Anna R. Karlin, and Elias Koutsoupias. The FedEx Problem. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, EC '16, pages 21–22, New York, NY, USA, 2016. ACM.

[19] Kira Goldner, Nicole Immorlica, and Brendan Lucier. Reducing inefficiency in carbon auctions with imperfect competition. In Thomas Vidick, editor, *11th Innovations in Theoretical Computer Science Conference, ITCS 2020, January 12-14, 2020, Seattle, Washington, USA*, volume 151 of *LIPIcs*, pages 15:1–15:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.

[20] Kira Goldner and Anna R. Karlin. A prior-independent revenue-maximizing auction for multiple additive bidders. WINE 2016, pages 160–173, Berlin, Heidelberg, 2016. Springer-Verlag.

[21] Jason D. Hartline and Tim Roughgarden. Optimal mechanism design and money burning. In *STOC '08: Proceedings of the 40th annual ACM symposium on Theory of computing*, pages 75–84, New York, NY, USA, 2008. ACM.

[22] Nicole Immorlica, Jacob Leshno, Irene Lo, and Brendan Lucier. Information acquisition in matching markets: The role of price discovery. *Available at SSRN*, 2020.

[23] Tony Ke and K. Sudhir. Privacy Rights and Data Security: GDPR and Personal Data Driven Markets. WINE 2020, 2020.

[24] Daniel Lehmann, Liadan Ita O'Callaghan, and Yoav Shoham. Truth revelation in approximately efficient combinatorial auctions. *Journal of the ACM (JACM)*, 49(5):577–602, 2002.

[25] Alejandro M. Manelli and Daniel R. Vincent. Multidimensional mechanism design: Revenue maximization and the multiple-good monopoly. *Journal of Economic Theory*, 137(1):153 – 185, 2007.

[26] R. Preston McAfee. A dominant strategy double auction. *Journal of Economic Theory*, 56(2):434–450, 1992.

[27] Paul R Milgrom and Robert J Weber. A theory of auctions and competitive bidding. *Econometrica: Journal of the Econometric Society*, pages 1089–1122, 1982.

[28] Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

[29] Roger B Myerson and Mark A Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of economic theory*, 29(2):265–281, 1983.

[30] Raghuvansh R. Saxena, Ariel Schvartzman, and S. Matthew Weinberg. The menu-complexity of one-and-a-half dimensional mechanism design. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '18, 2018.

[31] S. Matthew Weinberg and Zixin Zhou. Optimal multi-dimensional mechanisms are not local, 2020.

[32] Martin L Weitzman. Optimal search for the best alternative. *Econometrica: Journal of the Econometric Society*, pages 641–654, 1979.