



Minerva Access is the Institutional Repository of The University of Melbourne

Author/s:

Viljoen, S;Goldenfein, J;McGuigan, L

Title:

Design choices: Mechanism design and platform capitalism

Date:

2021-01-01

Citation:

Viljoen, S., Goldenfein, J. & McGuigan, L. (2021). Design choices: Mechanism design and platform capitalism. *Big Data and Society*, 8 (2), <https://doi.org/10.1177/205395172111034312>.

Persistent Link:

<https://hdl.handle.net/11343/289536>

License:

[CC BY NC ND](#)

Design choices: Mechanism design and platform capitalism

Big Data & Society
 July–December: 1–13
 © The Author(s) 2021
 Article reuse guidelines:
sagepub.com/journals-permissions
 DOI: 10.1177/20539517211034312
journals.sagepub.com/home/bds



Salomé Viljoen¹, Jake Goldenfein²  and Lee McGuigan³ 

Abstract

Mechanism design is a form of optimization developed in economic theory. It casts economists as institutional engineers, choosing an outcome and then arranging a set of market rules and conditions to achieve it. The toolkit from mechanism design is widely used in economics, policymaking, and now in building and managing online environments. Mechanism design has become one of the most pervasive yet inconspicuous influences on the digital mediation of social life. Its optimizing schemes structure online advertising markets and other multi-sided platform businesses. Whatever normative rationales mechanism design might draw on in its economic origins, as its influence has grown and its applications have become more computational, we suggest those justifications for using mechanism design to orchestrate and optimize human interaction are losing traction. In this article, we ask what ideological work mechanism design is doing in economics, computer science, and its applications to the governance of digital platforms. Observing mechanism design in action in algorithmic environments, we argue it has become a tool for producing information domination, distributing social costs in ways that benefit designers, and controlling and coordinating participants in multi-sided platforms.

Keywords

Mechanism design, digital platforms, optimization, programmatic advertising, labor platforms, algorithms

Introduction

Digital platforms have become infrastructures for organizing an increasing range of social, economic, and cultural activities (Caplan and Boyd, 2018; Nieborg and Poell, 2018; Plantin et al., 2018; Srnicek, 2017). Consternation about the power of these platforms is a hallmark of the present moment. Through automated and algorithmic tools and methods, companies like Google, Facebook, Amazon, and Uber exercise a special capacity for managing their customers, competitors, and workers (Darmody and Zwick, 2020; Shapiro, 2020; Yeung, 2017). They command many sources of leverage, including the monopoly power permitted under moribund antitrust enforcement (Khan, 2017; Srinivasan, 2020), regulatory blackholes that relieve them of social responsibilities (Rosenblat and Stark, 2016; Shapiro, 2018), and informational advantages owing to their position between buyers and sellers (Mansell and Steinmueller, 2020). To these concerns, we add a critical focus on the intellectual technologies powering these platforms—a specific set of disciplinary tools and methods deployed in their theorization, development, and operation. Core operations at Google, Facebook, and beyond derive from a research discipline that treats asymmetrical accumulation of data and profiling capacity, and the optimization of

actors' choices, as fundamental design principles. In this paper, we pull back the curtain of platform capitalism to glimpse the world of *mechanism design*.

Mechanism design brings a formidable but conflicted lineage into the heart of platform governance. With its origins in Nobel Prize-winning economic theory, mechanism design claims to be a set of methods tailored expressly to achieve social welfare by harnessing the self-interested rationality and autonomy of individuals. Mechanism design creates bespoke markets and auctions that direct individuals' choices toward outcomes that maximize a formally defined social welfare. Its ideological and normative commitments to rationality and autonomy have provided market-makers with a persuasive justification for exercising strategic control over the environments where people interact and make decisions. The apparent power of mechanism

¹Columbia University, USA

²Melbourne Law School, Australia

³University of North Carolina at Chapel Hill, USA

Corresponding author:

Lee McGuigan, University of North Carolina at Chapel Hill, Chapel Hill, USA.

Email: leemcg@email.unc.edu



design to orchestrate forms of social coordination has earned multiple Nobel Prizes for the discipline's leading proponents and helped to propagate its influence beyond conventional economic settings. Mechanism design is now regarded as an authoritative and generalizable science for engineering choice and distributing value in society (Maskin, 2008). It is seen to provide a set of intellectual technologies for enacting "optimization" as a worldview and a powerful idiom for a great deal of policymaking. And it has gained particular currency in areas where computers and algorithms mediate the actions of agents (Einav and Levin, 2014; Papadimitriou, 2001; Varian, 1995, 2010).

Largely without critical scrutiny, computational adaptations of mechanism design (and related management sciences, such as operations research) have become a major force in governing online environments and platforms. This so-called algorithmic or automated mechanism design (AMD) is now one of the most pervasive yet inconspicuous influences on the digital mediation of social life. The experiences of almost all web and mobile app users are organized, in part, by optimizing schemes from mechanism design—such as the advertising auctions that monetize online traffic and the decision systems that configure Facebook newsfeeds. As Google's chief economist Hal Varian (2010: 4) puts it, "Online advertising serves as a poster child for algorithmic mechanism design." There is accordingly considerable scope for augmenting discussions of platform and algorithmic governance with an analysis of how the intellectual tools of mechanism design are applied in those environments.

We argue that these computational applications warrant a critical reassessment of how this discipline tries to design choices. Through case studies in platform contexts, we show how mechanism design becomes a normative rhetoric of market facilitation that instead enables forms of market simulation designed to benefit the mechanism designer. Contrary to the legitimizing premises introduced above, the mechanism design methods underpinning digital platforms actually threaten social welfare and distort collective interests for platform firms' own ends. Platforms now use mechanism design to leverage data science and automation toward goals that dispense with the normative commitments assumed in its economic theorization. Drawing examples from online advertising auctions and the multi-sided market platforms that coordinate gig-labor and on-demand services, we illustrate how firms instrumentalize mechanism design to achieve information asymmetry, to distribute social costs in ways that benefit designers, and to orchestrate behaviors and choices in their systems. In light of the persistent accolades heaped on mechanism design, its enthusiastic adoption into platform settings exposes a contradiction between its theoretical premises and the industrial purposes that have operationalized them; a contradiction that has long defined the brackish mixture of mechanism design as theory and mechanism design as practice.

The paper offers four contributions to the scholarship on platform governance. First, it provides a critical introduction to mechanism design (and so-called AMD) for researchers who study big data, platform capitalism, and the values designed into technologies like algorithms and artificial intelligence. So far, most critical work in these areas has not engaged directly with this hybrid economic-engineering discipline, even though it is now central to the operations of platform firms, as well as other online networks and data-driven decision-making systems. Computer scientists are also turning to mechanism design as a framework for giving artificial intelligence a progressive social mission (Abebe and Goldner, 2018). A fluency in the vocabulary and techniques of mechanism design will help critical researchers translate and intervene in debates across these fields of science, policymaking, and industrial practice.

Second, we identify contradictions in the normative, ideological, and political commitments claimed by mechanism design's two main traditions: one descended from game theory and welfare economics, the other oriented around experimentation and applied computation. The digital uptake of mechanism design has papered over profound differences in how these two traditions define and operationalize key concepts, including the actor/subject, and its relationship to autonomous, rational decision-making.

Third, we look at how platform companies deploy and exploit mechanism design. We show how automated or algorithmic mechanism design operationalizes and progresses through the contradiction noted above, leveraging both the prestige and elegance of mechanism design's game theory tradition, as well as the brute force of the computational and statistical techniques honed by mechanism design's experimental tradition. Under the legitimizing glow of mechanism design, companies like Google and Facebook have optimized their platforms to accumulate information, capital, and managerial control.

Finally, we provide insights into how AMD extends the political project of installing automated *market-like* institutions to orchestrate social choices and allocate opportunities and value across more and more settings, while at the same time abandoning the normative commitment to autonomy and the ideological commitment to rationality that have been key to mechanism design's justification (Tomasetti, 2016). AMD is an example of what Mirowski and Nik-Khah (2017) call the "de-humanizing" of markets, in which, through their information-processing function, markets eschew any interest in human cognition or preferences. Put another way, the deployment of mechanism design in digital platforms carries forward the normative justification of markets into settings that appear like markets but operate more like control infrastructures.

Overall, our critique of mechanism design makes interventions in two directions: it gives critical researchers new insights into the discipline behind certain pathologies of platform capitalism; and it asks mechanism designers

to look critically at the contradictions and pathologies within the discipline itself.

Introduction to mechanism design

Mechanism design is a form of optimization developed in economic theory. It reverses the method of most economics. Typically, economists try to predict the outcomes of existing institutions. By contrast, mechanism design casts economists as institutional engineers who choose a desired outcome and then arrange a set of market rules and conditions to achieve this outcome. These rules and conditions (i.e. mechanisms) are formalized as algorithms—that is, prescriptive notations that can be demonstrated mathematically to optimize a selected outcome. The main thrust of mechanism design is to develop systems of social decision-making that direct actors' rational choices toward prescribed outcomes under conditions of uncertainty about actors' preferences or likely actions. Designers start with a goal, which could be to maximize efficiency, revenue, social welfare, or some other desideratum, and then they solve backwards, staging interactions expected to channel behaviors toward that goal.

Importantly, mechanism design is useful when the designer lacks (reliable) information that would allow an optimal allocation to be specified in advance. Auction mechanisms have been particularly celebrated for their ability to design around such information deficiencies. Some of the most significant work in mechanism design includes the design of auction formats and systems for matching scarce resources or opportunities with agents who want them (McAfee and McMillan, 1987; Myerson, 1981; Vickrey, 1961). These interests persist in work on AMD (Korula et al., 2016; McAfee and Vassilvitskii, 2012; Milgrom, 2019; Varian, 2009, 2010), where the ability to surface and leverage the private information held by individuals takes on even more significance for platform businesses organized largely to produce and process data rather than produce goods in the industrial sense.

Auction mechanisms

A brief introduction to auctions will help animate the normative commitments of mechanism design, elaborated upon below. Suppose a seller wants to sell an item to the prospective buyer who values it most—that is, has the highest willingness to pay. But the seller does not know how potential buyers value the item, and who among potential buyers values the item most highly; that is, the seller does not have access to the buyers' private information. The seller could simply set a price based on a guess about buyers' preferences, but that would risk pricing it below what someone would pay, or above what anyone would pay. Alternatively, the seller could ask buyers to disclose their values in a negotiation, but the former would have no reason to trust that the latter would report their preferences truthfully. So, instead, a

seller can use an auction to reveal that private information. The idea is to implement incentives and constraints that induce self-interested agents to express their true preferences (i.e. bid the amount that they actually value the item), and align their individual strategies and actions to the designer's optimal outcome (Clark, 1971; Vickrey, 1961).

Some of the auction types most widely adapted by platform businesses were theorized by William Vickrey, a Nobel Prize-winning economist. Vickrey's name is attached to a sealed-bid second-price auction for a single indivisible item. In this format, potential buyers each submit a bid privately; the highest bidder wins, but pays the price bid by the second highest bidder (or, often, a nominal increment above that "second price"). Vickrey mathematically showed that the sealed-bid second-price auction gives bidders incentives to bid their true values—to report honestly how much the item for sale is worth to them. The Vickrey auction has been, until recently, the prevailing format used in online display advertising. It was also adapted to apply to auctions that award multiple items to multiple bidders, as opposed to a single item. This creates a ranking and matching problem where optimization means maximizing the total value allocated among participants in the auction. This Vickrey–Clarke–Groves (VCG) auction also incentivizes truthful bidding as the optimal strategy, and it formulates the prices buyers pay as the cost inflicted by each bidder onto other bidders—by considering, for example, the value the second- and third-highest bidders would have claimed if the highest bidder had not participated in the auction. The VCG mechanism has been further generalized as a tool for determining socially-optimal outcomes in all sorts of matching or ranking problems. As discussed below, a convenient adaptation of the VCG, known as the generalized second price (GSP) auction, is the mechanism that conducts the enormously lucrative business of online search advertising (Edelman et al., 2007). Facebook also uses VCG auctions to place advertisements in users' news feeds (Sodomka, 2015).

There are many other auction mechanisms that can serve a variety of purposes. Despite these varying designs and applications, the normative premise is consistent: finding ways to channel private decisions based on privately held information into socially efficient distributions. This commitment follows from a particular economic lineage that lauds the market as a tool for social coordination, and it has helped justify mechanism design as a legitimate, defensible technique for managing the environments where individual actors make choices.

Mechanism design's origins, commitments, and justifications for orchestrating choices

Mechanism design has its normative origins in social choice theory, which animates the purposes, principles, and

assumptions of welfare economics. Social choice theory is a framework that studies how individual opinions, preferences, interests, or welfare can be combined in order to achieve some collective decision or aggregate social welfare (Arrow, 1963; Sen, 1999). Mechanism design evolved as one method with which to interrogate these problems of social choice by using applications from game theory to study “games” where agents have private information about their own preferences but may be strategically motivated to lie (Harsanyi, 1968).

Several insights from social choice theory and welfare economics provide the foundations of economic mechanism design. First, mechanism design relies on the general welfare economic axiom that if the correct conditions are achieved, a market allocation mechanism will lead to the socially optimal outcome of whatever good is at issue. This axiom derives from the utilitarian assumptions that (a) individuals are rational self-interested agents working to maximize their own welfare (as suggested by rational choice theory) and (b) that overall social welfare is simply the summation of each individual agent’s welfare. Therefore, the right conditions can guide individual action to achieve the maximum amount of total (social) welfare. This forms the basis of social welfare optimization as both a practice and method in economic theory. Second, mechanism designers operate under the assumption that under the right conditions, agents can and will express their true preferences—linking the expression of choice with that of autonomy in economic mechanism design. Together, this second assumption expresses what mechanism designers refer to as the “revelation principle” and the “implementability principle” (Vazirani et al., 2007). The revelation principle speaks to the methodological value of mechanisms: they don’t just record the values agents *say* they have, they construct interactions where agents show—reveal—the valuations they truly have. The implementability principle speaks to the methodological limits of mechanisms: this brand of truthful revelation can only be achieved under the right conditions. The task of a successful mechanism designer is to find and implement such conditions. The Vickrey auction described above satisfies these principles—in a true Vickrey auction, there is no benefit gained from bidding below one’s true valuation. The resulting mechanism will motivate agents to truthfully report their preferences and thus individually each pick the outcomes that collectively result in the desired social choice—a condition called the incentive compatibility constraint (Hurwicz, 1973). Setting this constraint mathematically and working from this fixed condition back to the settings required to achieve it is fundamental to the project of mechanism design.

These theoretical commitments are operationalized in mechanism design using agent-based methods from game theory. In game theory, agents may be apportioned a known degree of rationality, which they can then leverage

to act in games and learn about other players’ preferences, as well as their own, by incorporating information revealed through the game into their strategies (Bicchieri, 2004). In line with the general axioms laid out above, each rational agent is presumed to act, autonomously, based on its private information and preferences, in a sufficiently consistent way to maximize its own utility.

The autonomy-protecting elements of mechanism design make it an attractive bulwark against the autonomy-undermining and manipulative dimensions of digital technologies, often described as channeling informational inputs into desired behavioral outputs by creating and exploiting irrationality (Calo, 2014; Susser et al., 2019). These accounts diagnose the harm of behavioral manipulation as the undermining of autonomy, and the instrumentalization of users as behaviorist information processors.

Mechanism design, while also a form of choice architecture, does not intrinsically offend humanist conceptions of autonomy and rationality. Instead, it preserves a property- and market-centered vision of human autonomy, rationality, freedom, and subjectivity, while at the same time making certain presumptions about agents sharing fundamental characteristics and interests, which can be formalized into mathematical accounts of “rationality” (Mirowski and Nik-Khah, 2017).

This justificatory basis of mechanism design is expressed in the work of Eric Maskin (2015), who argues that mechanism design vindicates Friedrich Hayek’s account of markets as superlative information processing machines. Specifically, Maskin has argued that mechanism design supports the Hayekian hypothesis that the free market is preferable over central planning because of its information efficiency (including the claim that Pareto optimal mechanisms—where no actor can be better off without making at least one other actor worse off—would require actors to comprehensively report their preferences), and its enabling distributed actors to rationally use their independently held knowledge for self-interested purposes.¹

These origins in social choice and game theories have endowed mechanism design with a normative commitment to preserving actor autonomy and an ideological commitment to economic rationality. Mechanism design provides a framework, an algorithmic toolkit, for channeling autonomous and rational choices toward the social goals specified by designers.

Contradictions in mechanism design

But the full story of mechanism design is messier. Apart from concerns about both the privatizing thrust of treating all resources or opportunities as property, and the limits of determining value exclusively via utility, economic historians suggest that the realities of mechanism design in action deviate from its supposed theoretical generalizability. Furthermore, critics argue those general theories are

often fashioned to serve powerful (corporate) interests as much as to implement good policy. Mirowski and Nik-Khah (2017) argue that the auction format used by the US government to sell electromagnetic spectrum licenses in 1994 was not selected because of the undeniable elegance of game theory, as reported; behind the scenes, telecommunications companies lobbied for favorable game-theoretical designs, while experimental economists worked to adapt auction theory to the materialities of a workable technical system for administering the auctions in practice.

As Mirowski and Nik-Khah (2017) point out, mechanism design has always straddled these two discernible but convergent traditions—one rooted in game theory and the other in experimental methods. These traditions claim different normative commitments and assumptions about subjectivity and rationality, but they work from a generally shared ideology (belief in markets as the superlative technology for information processing and coordination of social choices) and a shared political project (of substituting markets for planning and bureaucracy). In contrast to the game theory tradition described above, the experimental approach to mechanism design is not concerned with preserving autonomous rational action, but rather with iterating toward the discovery of favorable correlations or statistical patterns between choice architectures and prescribed outcomes. Despite these different assumptions and commitments, these two traditions have leveraged each other's claims to authority, in complementary ways, to advance an ideological and political project that inheres in both economic mechanism design and AMD, and that has achieved a new crescendo with digital platforms. This distinguishes our critique of mechanism design from well-known criticisms of rational choice theory that challenge its formalization of human rationality on psychological (Sent, 2018), empirical (Green and Shapiro, 1994), sociological (Bourdieu, 2005), or ethical grounds (Hollis and Nell, 1975). Rather, we describe the willingness of mechanism design practitioners, where convenient, to undermine or even abandon mechanism design's own assumptions about how agents operate.

As mechanism design has travelled from economics into its second life online, the intrinsic contradictions between its stated goals and its effects in the world have amplified. This may reflect the practical conditions of its applications online (e.g. concentrated market power, infrastructures enabling massive data collection and processing, and corporate profit motives), but it also suggests mechanism design (as a design principle) demands a different kind of work than the collectively beneficial social ordering assumed as part of its early theoretical justifications. As more computational decision-making is integrated with inferential and predictive outputs from data science, mechanism design has become fundamental to processes that not only work to reveal preferences (according to the revelation principle), but also *infer actor preferences* from behavior or

associations through automated and recursively generated models. That is, rather than accepting that agents acting in a system possess private information inaccessible to the designers, those designers often leverage modular digital platforms to opaquely generate and exploit information asymmetries. These asymmetries help designers recognize the value created through agents interacting and transacting, and extract a greater share of surplus relative to the actors whose choices these designers try to conduct.

A second contradiction inherent to mechanism design relates to how rationality is operationalized. In economics, mechanism design developed as a way to enhance social welfare through individuals making decisions with their own private information and for their own desires and ends. Economic theory considers this morally and socially beneficial because it provides a way to hold individuals responsible for the social costs they generate through their personal choices, it enables individuals to coordinate free from coercion, and because the aggregate effects of such choices result in allocative efficiency from the markets they act in. In its algorithmic and computational applications, however, mechanism design evolved into a set of tools for instrumenting agent interactions under conditions of radical informational and calculative asymmetries via brute force experimentation (Shapiro, 2020). This leverages the social costs of decisions and actions in ways that benefit the designer. When deployed in concert with data science, mechanism design appears to withdraw any commitment to individual autonomy or the modelling of human cognition or rationality. Instead, human rationality becomes a *variable* to be tuned as designers search for desirable patterns of actions and rewards—a dramatic departure from the protection of human rationality that ostensibly legitimizes choice engineering as a means to execute social policy and governance.

Using platformed applications of mechanism design as case studies, below we describe how the justifications for mechanism design—both its respect for autonomy and rationality and its capacity to solve social coordination problems—are leveraged to legitimize a range of practices that actively negate the normative account mechanism design provides of itself. While these theoretical foundations may no longer cohere in mechanism design's automated or algorithmic applications, the ideological project of using “market” design to coordinate ever more spheres of human social activity persists. At the same time, we show how mechanism design participates in the contradictory process of automating away markets (Birch, 2020), by diminishing cognitive assumptions about agents, and moving information processing deeper into the algorithmic elements of the person-machine systems that constitute automated mechanisms. Put another way, as agent interactions are pushed further into privately controlled and highly automated digital platforms, agents are increasingly coordinated through algorithmic *market-like* mechanisms that

simulate how a market might behave without necessarily including any of the features necessary to constitute a market, such as freedom to deal or knowable information rules (Tomasetti, 2016). Instead, these market-like mechanisms are fully internalized by a single firm, and these mechanisms are characterized by information-rich, automated systems of iterative tuning.

Data science and AMD

When encountering computation and automation, the justificatory rationales described above do not automatically evaporate. For instance, the toolset of “algorithmic mechanism design” was leveraged to address several of the massive coordination problems surfaced by the development of large computer networks like the internet (Athens, 2010). These included networks operating “through the cooperation and competition of many participants, leading inevitably to underlying social and economic issues” (Blume et al., 2015, 2), and complex distributed systems where each agent is controlled by a different entity with its own independent goals that are not necessarily known to the protocol designer. In such situations, much like with the social choice theorists, the designer’s objective is to produce incentive compatible mechanisms (that promote truthfulness) to ensure rational systems or users who are only seeking to improve their own happiness comply with an online governing protocol (Nisan and Ronen, 2001). Classic mechanism design applications in the management of computer systems include SPAM reporting, dealing with freeloading in peer-to-peer file sharing, directing the movement of packets of information, and sharing computational resources.

But not all computational mechanism design applications operate under the archetypal information dynamic of unknown private information. Where continuing collection of information is possible (something digital technologies clearly afford, and economic logics clearly encourage [Sadowski, 2019; Srnicek, 2017; West, 2019]), mechanism design has been deployed to take advantage of those information dynamics, which greatly alters the capacities of the mechanism designer in relation to the agents interacting within the mechanism, and shifts the technique’s relationship to some of its theoretical and normative assumptions.

The “manual” form of mechanism design from economic theory requires intuition about the suitability of a rule set, and it provides formalisms for individual rationality and how the agents will interact under certain constraints. It does not rely on prior-held information about agents’ preferences. Where computational environments enable continuing collection of data about user behavior, and ongoing updating and tweaking of mechanisms—such as with AMD—agents’ preferences can be inferred using externally provided profiles of the agents, or by inference from behaviors expressed in multiple iterations of the mechanism in action. Some computational applications

design their mechanisms specifically to improve the inferability of preferences from actions. Certain platform applications described below will sacrifice revenue-optimal outcomes in order to obtain better data and understanding of the relationship between profile, behavior, and preference. Data about preferences becomes the primary goal, because of the long-term advantages afforded by information dominance in “platform” settings (Langley and Leyshon, 2017). Inferability thus becomes a new optimization target and mechanism design becomes a profiling tool.

AMD therefore specifies and reiterates its mechanisms according to the information that the designer is able to generate about the agents’ preferences (Sandholm, 2003). This is considered essential because static analytical mechanisms are presumed to be too clunky for dynamic real world uses like online advertising, airline pricing, ride sharing, and cloud server resource distribution. Mechanisms deployed in those environments require historical information to learn and optimize for future action (Niazadeh, 2017). To that end, some practitioners of AMD consider the perseverance of canonical mechanisms “astonishing” because they ignore the information available about agents in the system (Sandholm, 2003). Instead, automated mechanism designers infer agents’ types from actions (like bids), map actions to preferences through correlation, and then solve the mechanism to achieve whatever outcome is specified.

There are other meaningful differences between applied AMD and its theoretical progenitor. Where classical economics explores, for instance, single-item auctions, computational auctions occur billions of times each moment. That affords the computer scientist an opportunity to constantly adjust mechanisms with A/B testing, and to continually reiterate the mechanism through “mechanism responsiveness.” As the distribution of agent preferences evolves, so does equilibrium behavior, meaning the designer’s job is to continually observe and adapt the mechanism (Chawla et al., 2014).

Further, the statistical methods used in automated and computational approaches mark an important difference from mechanism design in more traditional economics, but are not *external* to mechanism design as a discipline. As mentioned above, economic historians have identified two different schools of mechanism design, even in its early developments (Nik-Khah and Mirowski, 2019). The first was the game-theoretic “Bayes-Nash school,” which produced generalizable theoretical formalizations of economic mechanism design, looking for socially beneficial forms of social coordination, by modelling interacting Bayesian agents. The second school, “the experimentalists,” was more interested in computer simulation and held different assumptions about the nature of markets and agents within them. Arguably, the latter school has had more practical influence in actual mechanism design applications, and appears to have supplied more of the intellectual and disciplinary tools and commitments for AMD.

Whereas the more analytical “Bayes-Nash” school used statistics to understand the influence of a few variables believed to have causal power, the experimentalists, especially as they began to deploy data science tools, used correlation to identify apparently meaningful relationships in high-dimensional datasets, along with automated model selection to identify the most meaningful or predictive variables (Einav and Levin, 2014). They were able to produce fully computational decision technologies that could be tweaked and reiterated through experimentation.

Given its origins in experimentalist mechanism design, AMD is less concerned with achieving “truthful” mechanisms or putting autonomous rationality to work. Automated mechanism designers have little interest in the quality of agents’ rationality. Far more important are the correlations between agent actions, preferences, and mechanisms. The agents operationalized in these mechanism systems are not the game theoretical rational agents of analytical mechanism design, but merely an empirical relation between model, action, and preference—a black-box linking informational inputs and behavioral outputs. It is frequently this kind of agent, whose behavior and decision making can be channeled algorithmically through massive data collection and repeated experimentation, that is put to work in the privately controlled market-like mechanisms of digital platforms.

Mechanism design in digital platform design, interaction, and governance

Platforms are instrumented to maximize the amount of data they collect from users, and through aggregation with data from other users, to process that data into predictions about users’ preferences (private information). This information dynamic and the asymmetry it puts to work, makes digital platforms rich terrain for amplifying the experimentalist tendencies of AMD. Search engines, social networks, and multisided platforms (MSPs) stage encounters with AMD in the daily lives of billions of people worldwide. Efforts to optimize advertising markets, for example, subtly but pervasively structure environments where people socialize, learn about the world, and find what they need to reproduce their bodies, identities, and relationships (Noble, 2018; Pasquale, 2015; Vaidhyanathan, 2018). Online advertising has been a catalyst for extensive theorizing and research in mechanism design, particularly on auctions (Milgrom, 2019). These designs leverage the authoritative force of the game-theory tradition, while following experimentalists’ techniques of modelling and data-driven iteration.

Most digital advertising is bought and sold using auction-based exchange platforms that “combine big data predictive models with sophisticated economic market mechanisms” (Einav and Levin, 2014: 12430895). As noted above, Google’s Hal Varian (2010: 4) has called

online advertising “a poster child for algorithmic mechanism design.” One of the strongest currents in digital marketing has been the disaggregation of audiences defined by contextual and demographic features into identifiable user-profiles classified according to potentially thousands of observed or inferred behaviors and attributes. Automated auctions are, in part, a mechanism for coping with advertisers’ desire to precisely and exclusively target valuable consumers—the “rational discrimination” (Gandy, 2009) afforded by combining surveillance, data science, and addressable communications. Rather than reserving placements on websites expected to deliver a bulk of user traffic, advertisers bid on discrete “impressions”—individual opportunities to serve ads—whose value they can better know or estimate. On the supply side of the market, auctions help publishers optimize their return on the inventory of audience attention they sell to advertisers. Auctions enable publishers to price and manage their audience inventory more rationally in a market of nearly limitless supply, and to better isolate the specific value of impressions associated with user profiles—to guess how much someone might be worth to an advertiser.

Both buyers and sellers want to profile web users in order to make more accurate guesses. But, as we will see, automated mechanism designers like Google—which administers auctions (and sometimes acts as buyer and seller all at once)—not only profile the users whose attention is for sale, but they also learn about the *bidders* in these auctions. Thus, in this section of the paper, the “actors” we focus on are not the people using the web and being exposed to ads, but the companies participating in online auctions. Using consumer profiles as inputs for price setting and adjudication of auctions, adtech intermediaries orchestrate the business-to-business transactions that finance media environments from which the behavioral responses of individuals can be procured—in a sort of *multi-sided mechanism design*.

Predictions and judgements in these markets flow at inhuman speed and scale. Researchers at Google point out that each impression is essentially unique, and so online display advertising involves the “daunting task” of pricing and selling “many trillions of distinct items” (McAfee and Vassilvitskii, 2012: 1059). Large advertising exchanges have been running tens of billions of auctions per day for most of the last decade (Korula et al., 2016: 30). In 2015, a Facebook researcher said the company was making trillions of decisions daily about how to price, rank, and deliver ads. He described the company’s advertising engine—which generated more than 98% of the \$70 billion in revenue Facebook reported in 2019—as being powered by an integration of machine learning and auction theory (Sodomka, 2015).

Algorithms and data science have become crucial to process trade at this volume and where the idiosyncrasies of each ad impression affect pricing and private value. The term “programmatic” advertising is generally used to

designate transactions involving auctions and/or computer-based automation and optimization (McGuigan, 2019; McStay, 2017). Its most notorious variant is “real-time bidding,” where an impression is auctioned and an ad placed in the 100–200 ms between when a user navigates to a webpage or opens a mobile app and when the content loads on the client device.

Participants across the programmatic market have turned to algorithmic mechanism design and machine learning to support adaptive, data-driven decisions. For instance, buyers try to refine their predictions about an impression’s value and competitors’ bids; sellers try to set an optimal reserve price, or floor, to maximize revenue; and platform intermediaries like Google and Facebook try to predict probabilities of clicks or other events following from a given ad, and they use those predictions to adjudicate bids. The dynamism of programmatic advertising also allows for sensitivities to information that would not be available or actionable via other transactional mechanisms. With location tracking and inferences about emotional or cognitive states, for example, advertising transactions can better account for the value to marketers of accessing individuals at a pivotal moment in their decision making and/or with insight into some discernible inclination (McStay, 2017). This information about the value of the inventory (i.e. the profile of the viewer and their likelihood of clicking on an ad) is then selectively managed through the ad mechanism via techniques that optimize price discrimination and bidding strategies in the ad action to squeeze value from these data and analytical resources. The belief that consumer profiling and predictive modeling provide credible means for making reliably (and quantifiably) better assessments of value has become a hallmark of this industry.

Mechanism design therefore anchors the business models of Google and Facebook, by far the biggest digital ad companies. Both use auctions and algorithms to sell and allocate ads. The design of these auctions has serious implications for the behaviors of the participants involved, and thus for the funding and organization of online information services. Pricing formulas, bidding sequences, access to data, and other variables can all impact the strategies of buyers and sellers. Algorithmic mechanism design in online advertising involves repeated interactions among participants, implying information gathering, experimental learning, and iterative optimization. Actions that may not be rational within an isolated auction event may pay off if a bidder or seller can discover over time a strategy that increases their share of the surplus between prices paid and buyers’ true values in subsequent auctions. Interest in these optimization models motivates lively exchanges in economics, operations research, data science, and among the adtech platforms providing bidding and/or inventory-management services.

Google is the biggest of those adtech providers. It operates the top programmatic exchange for display advertising, auctioning impressions for millions of publishers. More

importantly, its keyword search advertising has been phenomenally lucrative and deeply discussed in economics literature (Edelman et al., 2007; Varian, 2009). A typical keyword auction matches advertisers to several slots appearing around organic search returns and bills advertisers when a user clicks on an ad. To place ads that will generate revenue, Google assigns bidders a quality score that accounts for features expected to predict click-through-rates and user satisfaction (Varian, 2009). It then allocates ad positions using what it calls a GSP auction, where the advertiser pays the minimum price needed to preserve its rank among bidders, given the quality scores and expected click-through-rates (Edelman et al., 2007).

Google’s GSP format is a convenient adaptation of the quintessential VCG auction. A VCG auction internalizes to each bidder the social cost of their bid; its pricing formula charges individuals for the cost (or loss of value) their participation inflicts on other bidders. VCG therefore incentivizes truthful bids (satisfying the “revelation principle”), and it matches buyers to the items for sale in a socially optimal ranking. GSP simplifies these conditions, making the pricing formula easier to explain to advertisers. But the GSP sacrifices some of VCG’s socially beneficial characteristics, most notably those that secure truthful bidding and socially optimal equilibria (Edelman et al., 2007: 247). It is, therefore, serving different commitments than those for which VCG was theorized. Still, GSP has amassed a fortune for Google. And it improved on the first-price auctions that initially emerged in search advertising, which encouraged constant experimentation in advertisers’ strategies, including “socially inefficient investments into bidding robots” (Edelman et al., 2007: 245–246). Noting its achievements, one study calls GSP the “Cinderella of mechanism design” (Wilkens et al., 2017).

Google’s success testifies to the power of algorithmic mechanism design. But, as the GSP-VCG comparison suggests, its domination betrays practical realities and influences that often disrespect theoretical conditions and motivations relevant to mechanism design’s normative justification. Google operates a comprehensive adtech “stack,” providing technology, data, and logistical services to buyers, sellers, and intermediaries across advertising formats. Its position contributes to its market power and asymmetrical information advantages. For example, in search advertising, there is evidence that buyers in most Google auctions pay the reserve price (Scott Morton and Dinnielli, 2020: 13), suggesting that Google has learned to set price floors reliably between the highest and second highest bids. By optimizing the reserve price, Google is effectively bidding up its take by design. In the display advertising market that allocates ads and ad spending across the Web, Google’s near monopoly of the ad server and ad exchange businesses for years let the company have the first and last look at impressions and bids in programmatic auctions, amplifying its ability to recognize

value, price inventory, and issue bids itself (Srinivasan, 2020). Recent disclosures suggest that Google has been consciously leveraging this position as a multi-purpose platform intermediary to cultivate and exploit informational advantages, or what we might consider profiling capacities (Horwitz and Hagey, 2021).

Google's technical designs and business strategies defy assumptions about socially optimal mechanisms. The point of an auction is to incentivize truthful disclosures of preference and valuation; by contrast, Google, acting as an auction platform and multi-sided mechanism designer, is effectively profiling market actors so that it can extract more surplus for itself or its advertiser clients. As a recent report notes, "It is unusual, to say the least, for a single company to represent both sellers and buyers in the same market, and also to set the rules for, and conduct, the auctions that determine the winners, losers, and prices in that market" (Scott Morton and Dinnielli, 2020: 19; see also Srinivasan, 2020).

Facebook is another paragon of algorithmic mechanism design. The social network is programmed to optimize "engagement," which aims to operationalize value for all participants—meaning people see ads and organic posts that interest them, and advertisements are inserted into situations calculated as likely to meet advertisers' objectives. Facebook's advertising decision system uses machine learning to evaluate ads and user-generated content, and it applies a VCG mechanism to rank, configure, and price the ads on a page (Sodomka, 2015). Auctions determine winners by considering together the bid, a measure of the quality of the ad, and an estimate of whether the profiled user will take a desired action when exposed to the ad. In delegating these ad-serving decisions to Facebook, advertisers essentially bet on Facebook's powers of discrimination—its capacity to recognize and act on differences in value or probability.

Facebook and other digital advertising platforms, which seem perfectly suited for mechanism design, transport economic, and technical theories into situations where their elegance takes a back seat to brute force. Digital advertising is a relentless, iterative, and automated process of bidding, tweaking messages, learning probabilities, and updating predictions for the events advertisers can bid on, and deciding on placements to meet those goals with the optimal use of space. Facebook's platform embodies the spirit of mechanism design in that advertisers ostensibly pay for Facebook to engineer a selected behavioral *outcome* (e.g. "Likes," page visits). Importantly though, rather than staging choices aligned to agents' theorized rationality, Facebook experiments with computationally feasible designs to configure whatever expression of "rationality" yields satisfactory correlations. Describing the machine learning techniques Facebook has used to predict click rates and make ad-serving decisions, a researcher said, "We don't always have good theoretical bases for the

things that we try, but it's just so easy to try them that sometimes that has actually worked out pretty well" (Sodomka, 2015).

Mechanism design in MSPs

Social networks and online advertising exchanges are examples of the broader category of MSPs—businesses that match buyers and suppliers for a range of products and services. Many consumer-facing MSPs (eBay, Uber, Airbnb) have become well known and highly capitalized. Often presenting their product as market-making technology, MSP companies promise to reduce search and/or transaction costs for participants. But MSPs are not neutral intermediaries. Arguably, they are not even really markets. Instead, they are tools and enactments of world building; their designs and optimization strategies channel power, structure forms of decision making, and encode economic and legal relationships.

In these environments, mechanism design techniques are used to manage platform participants in ways that would otherwise implicate employment or antitrust law (Rosenblat and Stark, 2016; Shapiro, 2018). Consider the status of "gig workers" whom customers access via on-demand service apps, like Uber and DoorDash. On one hand, by granting workers autonomy in choosing when and how to work, MSP companies can designate workers as "independent contractors," thus avoiding costly obligations (e.g. employee benefits). On the other hand, worker autonomy is a barrier to controlling labor costs and rationally distributing workers geographically and temporally to serve customers on the other side of the platform.

Companies have tried to resolve this tension by using mechanism design to configure the settings in which choices are made, rather than by coercing such choices explicitly. Studies of ride-hailing and on-demand couriers demonstrate how platforms try to manage their workforces by cultivating informational and calculative asymmetries (Rosenblat and Stark, 2016; Shapiro, 2018, 2020). The companies attempt to engineer preferred choices among workers by manipulating the flow of information and, ultimately, configuring the decision-making resources available through platform interfaces. Like other mechanism designers, platforms establish the "calculative equipment" (Callon and Muniesa, 2005) available to guide choices; but platforms also exploit asymmetrical power over both market rules and market information. MSPs can analyze the interactions they process and then adapt choice architectures to take advantage of apparent patterns. Aaron Shapiro (2020) points to dynamic or surge pricing in on-demand apps as a paradigmatic example of an institution built both to capture more of the surplus between prices and customers' willingness to pay, and to adjust the distribution of available drivers. Price-setting, here, is not just a response to supply and demand signals, but an instrument

for influencing behavior on both sides of the platform toward managing the probabilities of possible outcomes, affecting how consumers and workers calculate their private values. Studies suggest surge pricing does not even reflect an empirical measurement of demand, but rather a prediction or pre-emption calculated to minimize rider waiting times by algorithmically managing drivers to “autonomously” shift into a particular geographic area (Rosenblat and Stark, 2016). The idea is not so much to harness actors’ rationality as to define and modulate their rationality—to model consumer preferences and workers’ responsiveness to changing payouts, for example, in ways that help platforms anticipate and select governance options (e.g. a mix of incentives/disincentives, pricing algorithms, etc.).

Both programmatic advertising and on-demand apps illustrate how online platforms use mechanism design selectively to formalize and implement optimization strategies that leverage informational, calculative, and power asymmetries. These applications of mechanism design to digital environments—particularly where computational systems, like bidding and pricing algorithms, interact with each other—are viewed optimistically as market builds that can transcend human capacities for information processing and rationality. In other words, algorithms and AIs, not humans, might best resemble the agents and decision makers posited in economic game theory (Varian, 1995). Because interactions among those automated rational agents can be programmed more easily than human-scale mechanisms, automated mechanisms become an attractive way to design new market rules, rendering market design (classically defended and theorized as an allocative mechanism with desirable political properties) “as much a matter of science and engineering as it is of public policy” (Parkes and Wellman, 2015: 272). But because digital platforms mediate densely woven economic, cultural, political, legal, and logistical relationships (Plantin et al., 2018), the most elegant theoretical solutions to social choice problems may not fit into the fabric of existing social contexts. According to theoretical computer scientist Christos Papadimitriou (2001: 752), designers must configure incentives that clear a “path” toward implementation. With the complex socio-economic entanglements of internet-enabled technologies, he says, “All design problems are now mechanism design problems.”

The politics of economic-oriented digital design

When accepting his Nobel Prize, Maskin (2008: 568) described the value of mechanism design as a tool to implement social goals. Speaking of public goods, he notes, “The net surplus-maximizing choice of public goods depends on citizens’ preferences for such goods, and there is no particular reason why the government should know these

preferences.” This understanding of both the intellectual challenge of mechanism design, and its supporting theory of a disinterested designer facilitating interactions between autonomous agents, strikes a discordant note with the platform designs discussed above. What to make of such an apparent disconnect?

One view may be that mechanism design in its grand, analytical, generalizable, game-theoretic, Nobel Prize-winning forms, and its applications in online platforms represent a form of disciplinary drift—an industrial and algorithmic perversion of a promising set of policy tools. Alternatively, it may be that mechanism design has always held this contradiction within itself, and leveraged the normative and ideological justifications of economic mechanism design to become a pervasive and generally accepted instrument for social coordination. Indeed, some of the particularities we describe in the context of online platforms can be clearly traced back to the experimentalists. Experimental mechanism design has origins in operations research, software simulation, public-choice economics, and the imagination of markets as hybrid person-machine systems capable of replacing bureaucratic governance. As the model of markets implemented in mechanisms has become more “algorithmic,” and the “markets” they engineer look more like automated decision systems, mechanism design may in effect eschew its commitment to optimum social allocations, while nonetheless furthering the ideological program of conceiving of markets as essentially and fundamentally information processors. Indeed, if mechanism design represents vindication of the idea that markets can be conceived as information processors, then it should not surprise us that the agendas those “markets” have come to serve are the agendas of the most powerful information processing firms in history.

The game theoretical models used in traditional economic mechanism design were based on the interest and utility maximizing (i.e. classical economic) versions of rationality (Weirich, 2004), considered fundamental to individual autonomy (Bicchieri, 2004). And while some computational approaches extend and apply these versions of rationality (Halpern and Pass, 2015), the data scientific and platform applications deployed in the platform space leave these accounts of rationality and causation behind. All that matters is statistical correlation between behavior and preference, and evolutionary convergence of a mechanism on its optimization target.

This operationalizes the behaviorist vision that sees individuals as black-boxed symbolic information processors responding to information inputs. As Nik-Khah and Mirowski (2019: 54) note, “market designers did not view the ‘person’ part of this ‘person-machine system’ with much in the way of cognitive capacity.” Instead, markets needed to be robust to the cognitive failures of persons, and operate without much concern for what individuals actually want. In the platform space where the data

scientific notions of rationality are put to work, whatever fundamental normative account that privileges mechanism design over other forms of social allocation and cooperative decision-making therefore loses traction. Mechanism designers infer preferences from behavior, and through profiling claim to have access to better information about agents' preferences than even those agents would be capable of reporting themselves. Similarly, in multi-sided platforms, large players have incredible quantities of information about the supply and demand data of producers. Further, while it may not be possible to design incentive compatible mechanisms, that does not seem to be the point of mechanism design as deployed in digital platforms anyway. Somewhere in the slippage out of economics and into computer science, the deployment of mechanism design to achieve social choice was marginalized in favor of using mechanism design to achieve the self-interested ends of the planner, and the tools for doing so made human rationality a variable. But this is also an inevitable consequence of a move that *automates away markets*, handing over control of that automation to firms capable of monopolistic information processing that are less interested in social coordination, but instead capturing surplus from the value created by agential interaction (Birch, 2020).

Indeed, while mechanism designers understand their task as the design and engineering of markets to provide bespoke information processing solutions to coordination problems, sociologists and economic historians have shown that these mechanisms do not really look—or operate—like markets at all. Instead, mechanisms are better understood as algorithmic decision systems, involving human and computational elements that circumvent the difficulties of modelling human cognition to produce outcomes in service of corporate and industrial interests. Contrary to mechanism design's premise of market freedom and autonomous actors, in practice the substitution of markets by private bureaucratic mechanisms of algorithmic management perhaps works to vindicate something else: theories of the behaviorist subject, for whom individually held preferences and choices break down in the face of recursively optimized mechanisms designed to take advantage of statistical behavioral response under conditions of massive information and calculative asymmetries.

Conclusion

Like social choice theory more broadly, mechanism design has gained rapid influence beyond economics, solving questions of institutional design across a number of policy settings, and influencing political science, philosophy, mathematics, and more recently computer science. As the method moved from theoretical inquiries of mathematical conditions required to coordinate individual action for the social good, to the real-world settings in which practitioners sought to realize those mathematical conditions,

these theories encountered both the functional reality of experimental and computational methods required to implement them, and the political reality of being taken up to further particular sets of interests. Together these serve to strain the relationship between the ideological and theoretical commitments that motivated mechanism design in theory, and systems developed from mechanism design in the wild. This contradiction or tension becomes particularly heightened in the application to digital platform settings, where the drift from a theoretical inquiry of market conditions into a mode of engineering decision technologies is particularly pronounced.

We see mechanism design as a normative tool for designing and building our world. Mechanism designers are engineers—world builders. There are implicit normative theories underlying all of mechanism design, and mechanism design becomes normative in its deployment to build economic institutions and tools. We see this normative justification across policy domains and embedded in the ideology of platforms as new types of participatory market environments. Understanding mechanism design in action, and the way in which it eats its own assumptions when moving across historical and disciplinary boundaries, therefore offers us a new story about the topography of digital pseudo-markets (i.e. platforms), their deployments of optimization and automation, and how market rationality is a contingent and constructed notion that defines us as economic and social actors for the benefit of the designers of the digital worlds we inhabit.

Acknowledgments

The authors thank the reviewers for their helpful comments. Much is owed to colleagues at Cornell Tech's Digital Life Initiative, as well as New York University Law School's Information Law Institute. Many thanks to Fenwick McKelvey and Josh Neves for their assistance with an earlier version of this article. The authors also thank Jan Wittington for insightful commentary when this work was presented at Privacy Law Scholars Conference 2020, as well as many others in attendance that offered suggestions including David Robinson. The authors also thank organizers and attendees at other meetings where this work was presented: the Mechanism Design for Social Good 2020 workshop; the Center for Information Technology and Public Life at UNC-CH; and the Center for Human-Compatible AI at UC Berkeley.


Declaration of conflicting interests


The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The authors received no financial support for the research, authorship and/or publication of this article.

ORCID iDs

Jake Goldenfein  <https://orcid.org/0000-0002-0242-8559>

Lee McGuigan  <https://orcid.org/0000-0003-3157-2944>

Note

1. Scholars have challenged the relationship between individual rationality and market performance that Maskin attributes to Hayek (Slobodian, 2018).

References

- Abebe R and Goldner K (2018) Mechanism design for social good. *AI Matters* 4(3): 27–34.
- Arrow KJ (1963) *Social Choice and Individual Values*. New Haven: Yale University Press.
- Athens G (2010) Mechanism design meets computer science. *Communications of the ACM* 53(8): 11–13.
- Bicchieri C (2004) Rationality and game theory. In Mele A and Rawling P (eds) *The Oxford Handbook on Rationality*. New York: Oxford University Press, pp. 182–205.
- Birch K (2020) Automated neoliberalism? The digital organisation of markets in technoscientific capitalism. *New Formations* 100–101: 10–27.
- Blume L, Easley D, Kleinberg J, et al. (2015) Introduction to computer science and economic theory. *Journal of Economic Theory* 156: 1–13.
- Bourdieu P (2005) *The Social Structures of the Economy*. Polity: Cambridge.
- Callon M and Muniesa F (2005) Peripheral vision: Economic markets as calculative collective devices. *Organization Studies* 26(8): 1229–1250.
- Calo R (2014) Digital market manipulation. *George Washington Law Review* 82: 995–1051.
- Caplan R and Boyd D (2018) Isomorphism through algorithms: Institutional dependencies in the case of Facebook. *Big Data & Society* (Jan–June): 1–12.
- Chawla S, Hartline J and Nikipelov D (2014) Mechanism design for data science. Proceedings of 15th ACM Conference on Economics and Computation, Palo Alto, CA, pp. 711–712.
- Clark EH (1971) Multipart pricing of public goods. *Public Choice* 11(1): 17–33.
- Darmody A and Zwick D (2020) Manipulate to empower: Hyper-relevance and the contradictions of marketing in the age of surveillance capitalism. *Big Data & Society* (Jan–June): 1–12.
- Edelman B, Ostrovsky M and Schwarz M (2007) Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review* 97(1): 242–259.
- Einav L and Levin J (2014) Economics in the age of big data. *Science* 246(6210), 12430891–12430896.
- Gandy OH Jr (2009) *Coming to Terms with Change: Engaging Rational Discrimination and Cumulative Disadvantage*. Farnham, UK: Ashgate Publishing.
- Green DP and Shapiro I (1994) *Pathologies of Rational Choice Theory*. New Haven: Yale University Press.
- Halpern JY and Pass R (2015) Algorithmic rationality: Game theory with costly computation. *Journal of Economic Theory* 156: 246–268.
- Harsanyi JC (1968) Games with incomplete information played by ‘Bayesian’ players, I–III Part I The basic model. *Management Science* 14(5): 320–334.
- Hollis M and Nell EJ (1975) *Rational Economic Man*. New York: Cambridge University Press.
- Horwitz J and Hagey K (2021) Secret Google Project Accused of Using Ad Data to Lift Sales. *Wall Street Journal* 277(84): A1, A9.
- Hurwicz L (1973) The design of mechanisms for resource allocations. *American Economic Review* 63(2): 1–30.
- Khan LM (2017) Amazon’s antitrust paradox. *Yale Law Journal* 126(3): 710–805.
- Korula N, Vahab M and Nazerzadeh H (2016) Optimizing display advertising markets: Challenges and directions. *IEEE Computing* 20(1): 28–35.
- Langley P and Leyshon A (2017) Platform capitalism: The intermediation and capitalisation of digital economic circulation. *Finance and Society* 3(1): 11–31.
- Mansell R and Steinmueller WE (2020) *Advanced Introduction to Platform Economics*. Cheltenham, UK: Edward Elgar Publishing.
- Maskin E (2008) Mechanism design: How to implement social goals. *American Economic Review* 98(3): 567–576.
- Maskin E (2015) Friedrich von Hayek and mechanism design. *The Review of Austrian Economics* 28(3): 247–252.
- McAfee RP and McMillan J (1987) Auctions and bidding. *Journal of Economic Literature* 25(2): 699–738.
- McAfee RP and Vassilvitskii S (2012) An overview of practical exchange design. *Current Science* 103(9): 1056–1063.
- McGuigan L (2019) Automating the audience commodity: The unacknowledged ancestry of programmatic advertising. *New Media & Society* 21(11–12): 2366–2385.
- McStay A (2017) Micro-moments, liquidity, intimacy and automation: Developments in programmatic ad-tech. In Siegert G, von Rimscha MB and Grubenmann S (eds) *Commercial Communication in the Digital Age – Information or Disinformation?* Munich: De Gruyter, pp. 143–159.
- Milgrom P (2019) Auction market design: Recent innovations. *Annual Review of Economics* 11: 383–405.
- Mirowski P and Nik-Khah E (2017) *The Knowledge We Have Lost in Information*. New York: Oxford University Press.
- Myerson RB (1981) Optimal auction design. *Mathematics of Operations Research* 6(1): 58–73.
- Niazadeh R (2017) Algorithms vs. Mechanisms: Mechanism Design for Complex Environments. PhD Thesis, Cornell University, Ithaca, NY.
- Nieborg DB and Poell T (2018) The platformization of cultural production: Theorizing the contingent cultural commodity. *New Media & Society* 20(11): 4275–4292.
- Nik-Khah E and Mirowski P (2019) The ghosts of Hayek in orthodox microeconomics: Markets as information processors. In Beverungen A, Mirowski P, Nik-Khah E and Schröter J (eds) *Markets*. Minneapolis: University of Minnesota Press, pp. 31–70.
- Nisan N and Ronen A (2001) Algorithmic mechanism design. *Games and Economic Behavior* 35(1–2): 166–196.
- Noble SU (2018) *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York: NYU Press.
- Papadimitriou CH (2001) Algorithms, games, and the Internet. *STOC ‘01*: 749–753, 752.
- Parkes DC and Wellman MP (2015) Economic reasoning and artificial intelligence. *Science* 349(6245): 267–272.

- Pasquale F (2015) *The Black Box Society: The Secret Algorithms That Control Money and Information*. Cambridge: Harvard University Press.
- Plantin J, Lagoze C, Edwards PN, et al. (2018) Infrastructure studies meet platform studies in the age of Google and Facebook. *New Media & Society* 20(1): 293–310.
- Rosenblat A and Stark L (2016) Algorithmic labor and information asymmetries: A case study of Uber's Drivers. *International Journal of Communication* 10(27): 10–27.
- Sadowski J (2019) When data is capital: Datafication, accumulation, and extraction. *Big Data & Society* (Jan–June), 1–12.
- Sandholm T (2003) Automated mechanism design: A new application area for search algorithms. International Conference on Principles and Practice of Constraint Programming (September), pp. 19–36.
- Scott Morton FM and Dinnelli DC (2020) *Roadmap for a Digital Advertising Monopolization Case Against Google*. Omidyar Network. <https://www.omidyar.com/sites/default/files/Roadmap%20for%20a%20Case%20Against%20Google.pdf>
- Sen A (1999) The possibility of social choice. *American Economic Review* 89(3): 349–378.
- Sent EM (2018) Rationality and bounded rationality: You can't have one without the other. *The European Journal of the History of Economic Thought* 25(6): 1370–1386.
- Shapiro A (2018) Between autonomy and control: Strategies of arbitrage in the 'on-demand' economy. *New Media & Society* 20(8): 2954–2971.
- Shapiro A (2020) Dynamic exploits: Calculative asymmetries in the on-demand economy. *New Technology, Work, and Employment* 35(2): 162–177.
- Slobodian Q (2018) *Globalists: The End of Empire and the Birth of Neoliberalism*. Cambridge: Harvard University Press.
- Sodomka E (2015) On how machine learning and auction theory power Facebook advertising. Simons Institute for the Theory of Computing, Berkeley, CA, November 17, 2015. Video, 54:21. <https://simons.berkeley.edu/talks/eric-sodomka-2015-11-17>
- Srinivasan D (2020) Why Google dominates advertising markets. *Stanford Technology Law Review* 24(1): 55–174.
- Srnicek N (2017) *Platform Capitalism*. Cambridge, UK: Polity.
- Susser D, Roessler B and Nissenbaum H (2019) Technology, autonomy, and manipulation. *Internet Policy Review* 8(2): 1–22.
- Tomasetti J (2016) Does Uber redefine the firm? The post-industrial corporation and advanced information technology. *Hofstra Labor & Employment Law Journal* 34(1): 1–78.
- Vaidhyanathan S (2018) *Antisocial Media: How Facebook Disconnects Us and Undermines Democracy*. New York: Oxford University Press.
- Varian HR (1995) Economic mechanism design for computerized agents. Proceedings of the First USENIX Workshop on Electronic Commerce. New York, NY, USA, pp. 13–21.
- Varian HR (2009) Online ad auctions. *American Economic Review* 99(2): 430–434.
- Varian HR (2010) Computer mediated transactions. *American Economic Review* 100(2): 1–10.
- Vazirani VV, Nisan N, Roughgarden T, et al. (2007) *Algorithmic Game Theory*. Cambridge: Cambridge University Press.
- Vickrey W (1961) Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance* 16(1): 8–37.
- Weirich P (2004) Economic rationality. In Danielson P, Audi R and Bicchieri C (eds) *The Oxford Handbook of Rationality*. New York: Oxford University Press, pp. 380–398.
- West SM (2019) Data capitalism: Redefining the logics of surveillance and privacy. *Business & Society* 58(1): 20–41.
- Wilkens CA, Cavallo R and Niazadeh R (2017) GSP: The Cinderella of mechanism design. Proceedings of the 26th International Conference of the World Wide Web, Perth, Australia, pp. 25–32.
- Yeung K (2017) 'Hypernudge': Big data as a mode of regulation by design. *Information, Communication & Society* 20(1): 118–136.