

The Parable of the Auctioneer: Complexity in Paul R. Milgrom’s *Discovering Prices*

Scott Duke Kominers* and Alexander Teytelboym**

Designing marketplaces in complex settings requires both novel economic theory and real-world engineering, often drawing upon ideas from fields such as computer science and operations research. In *Discovering Prices*, Milgrom (2017) explains the theory and design of the United States “Incentive Auction” that reallocated wireless spectrum licenses from television broadcasters to telecoms. Milgrom’s account teaches us how economic designers can grapple with complexity both in theory and in practice. Along the way, we come to understand several different types of complexity that can arise in market design.

Keywords: auctions, pricing, market design, complexity

JEL Classification: D47, C78, D44, D82, D02, D51, D71, D61, D62, D63

1. Introduction

For hundreds of years, people have found ingenious ways to allocate resources and discover prices. Merchants in medieval Europe, for example, set up multilateral clearing mechanisms for non-tradable debt, known as *rescontre*; these decentralized matching algorithms involved first clearing reciprocal debts and then clearing debt chains and cycles (e.g., in which i owes j who owes k who owes i). *Rescontre* appears to have been highly efficient (Börner and Hatfield, 2017).¹

Meanwhile, the Ming and Qing Dynasties (1594–1700) used randomized assignment systems to eliminate the possibility of corruption in the selection of candidates for civil service jobs (Will, 2002).² And even the great German poet Johann Wolfgang von Goethe designed an auction: He was interested in learning what value his publisher put on a manuscript, and this led him to devise a mechanism that turned out to be equivalent to a second-price, sealed-bid auction (Moldovanu and Tietzel, 1998).³

Although marketplace design has existed for ages, systematizing the field through the language

* Entrepreneurial Management Unit, Harvard Business School; Department of Economics, Center of Mathematical Sciences and Applications (CMSA), and Center for Research on Computation and Society, Harvard University; and National Bureau of Economic Research.

** Department of Economics, Institute for New Economic Thinking at the Oxford Martin School, and St. Catherine’s College, University of Oxford.

The authors appreciate helpful comments from the editor, Steven Durlauf, as well as from Zoë Hitzig, Paul Klemperer, Shengwu Li, Alex Nichifor, Yufei Zhao, and especially Paul R. Milgrom. Kominers is grateful for the support of the National Science Foundation (grants SciSIP-1535813 and SES-1459912), the Washington Center for Equitable Growth, the Ng Fund and the Mathematics in Economics Fund of the CMSA, and the Human Capital and Economic Opportunity Working Group (HCEO) sponsored by the Institute for New Economic Thinking (INET). Teytelboym is grateful for the support of the Economic and Social Research Council (grant ES/R007470/1). Both authors are grateful for the support of an Alfred P. Sloan Foundation grant to the CMSA.

¹In Venice, the *rescontre* mechanism was temporarily banned in 1593 (Börner and Hatfield, 2012, p. 15); this may have contributed to some moneylenders’ subsequent decision to demand corporal collateral, as in Shakespeare’s (1600) “pound of flesh.”

²Some have also speculated that Babylonians adopted lottery mechanisms for fairness reasons (Borges, 1959).

³Unfortunately, because Goethe’s legal counsel was unreliable, the publisher learned the reserve price in advance.

and logic of economic theory has allowed us to address problems in increasingly complex settings. The markets we consider today are more diverse and interconnected than before—and our design goals have grown more ambitious.

In the 1950s, the American Medical Association (AMA) faced the problem of fairly and efficiently allocating residents to hospitals. The AMA eventually arrived at an algorithm that appeared to work rather well—the elegant [Gale and Shapley \(1962\)](#) algorithm ([Roth, 1984, 2008](#)). But several decades later, the residency match faced a new problem: Doctors were marrying each other in greater numbers, necessitating a method to incorporate couples’ joint preferences into the residency assignment process; solving this real-world “couples matching” problem relied on a subtle combination of economic theory and operations research ([Roth and Peranson, 1999](#)). Likewise, economic theory has informed the adoption of [Myerson \(1981\)](#) optimal reserve prices in online ad auctions ([Ostrovsky and Schwarz, 2016](#)), as well as the eventual shift of ad auctions from practitioner-designed “generalized second price” auctions ([Aggarwal et al., 2006](#); [Edelman et al., 2007](#); [Varian, 2009](#); [Athey and Ellison, 2011](#)) to auction theory-inspired [Vickrey–Clarke–Groves](#) mechanisms ([Athey and Nekipelov, 2010](#); [Varian and Harris, 2014](#)). It is hard to imagine any of these innovations in marketplace design arising (or even being possible) without the groundwork laid by economic theory.

Nowadays, many markets are extraordinarily complex—both in theory and in practice. Consider the myriad difficulties a ride-sharing company such as Lyft or Uber faces in organizing its marketplace: The company must figure out where its customers want to go and how much they are willing to pay, while tracking drivers and their time-varying reserve prices for providing service. The company then needs to efficiently calculate journey prices and match drivers to riders. Overall, the company must raise revenue to cover its costs (and reward its shareholders) while encouraging drivers and customers to join and remain on the platform. And all of this has to happen “under the hood” in real time, and preferably before customers get tired of waiting and switch to another app.

It seems intuitive that matching of couples to hospitals and optimal reserve pricing is somehow more “complex” than Goethe’s manuscript allocation problem—and the problems that ride-sharing companies face seems more complex still. *But what does “complex” actually mean in this context?* Economic theory does not on its own give us a language for understanding complexity of market(place) design problems. Indeed, as [Paul R. Milgrom](#) argues in his new book, *Discovering Prices: Auction Design in Markets with Complex Constraints*:

“most of economic theory treats simple models of the economy in which prices alone can guide efficient economic decisions, but that conclusion rarely applies exactly in complex systems” ([Milgrom, 2017](#), p. 44).

Discovering Prices is based on the Arrow Lecture [Milgrom](#) delivered at Columbia University in November 2014. The introductory chapter is itself a mini-monograph accessible to non-specialists (including non-economists); there, [Milgrom](#) illustrates the challenges that arise in designing markets when there are complex constraints on the allocative process. [Milgrom](#) discusses three examples: farm consolidation after land lotteries in Georgia; determining airspace rights for commercial space and atmospheric travel; and the Federal Communications Commission’s (FCC’s) multibillion-dollar “Incentive Auction,” which reallocated wireless spectrum to repurpose it from broadcast television to telecom applications. In all three settings, there are significant complementarities and interdependencies between goods; as [Milgrom](#) shows, this makes it difficult to find price mechanisms that guide market participants towards efficient allocations. Any economic theorist could read *Discovering Prices* from Chapter 2, which opens with a masterful exposition of the close analytical relationships between existence and stability of competitive equilibria, tâtonnement in general equilibrium theory, deferred acceptance algorithms in matching markets, and auction design for indivisible goods. As [Milgrom](#) points out, “substitutability” of goods links all four settings. Indeed, “gross

substitutability” is crucial for uniqueness of competitive equilibrium and for tâtonnement in a classical general equilibrium setting (Arrow and Hurwicz, 1958), for existence of equilibrium in markets with indivisible goods (Kelso and Crawford, 1982; Milgrom and Strulovici, 2009; Gul and Stacchetti, 1999; Sun and Yang, 2006; Hatfield et al., 2013; Fleiner et al., forthcoming), for convergence of simple monotonic auctions (e.g., the deferred acceptance algorithm) to efficient outcomes (Gul and Stacchetti, 2000; Sun and Yang, 2009; Teytelboym, 2014), and for the numerous attractive properties of the Vickrey auction (Ausubel and Milgrom, 2006).

Milgrom’s main theoretical innovations are the introduction and development of “near-substitutability,” and associated approximate equilibria (inspired by the pioneering work of Dantzig (1957) and Korte and Hausmann (1978)), along with the analysis of the performance of the deferred acceptance auction without substitutability (Milgrom and Segal, forthcoming). The final two chapters connect the concepts of substitutes and near-substitutes to powerful insights from operations research, in order to explain the theoretical foundations behind the Incentive Auction, and to suggest future directions for the design of auctions under complex constraints.

In this essay, taking inspiration from Milgrom (2017), we offer a conceptual framework for thinking about complexity in markets. We show how Milgrom’s account of the design of the Incentive Auction serves as a sort of “parable” for the design of complex markets. Along the way, we highlight three different types of complexity and describe how they shape marketplace design.

2. Markets, Marketplaces, and Complexity

A *market* arises whenever agents interact and seek to transact, subject to economic incentives. Markets are everywhere—and many (but not all) of them are at least partially organized through *marketplaces* consisting of *rules* guiding which transactions occur and how, and/or *infrastructure* that facilitates transaction execution (Kominers et al., 2017; Eisenmann and Kominers, 2018; Kominers, 2018; Roth, 2018).

Marketplace design operates in *settings* consisting of

- (1) a market and
- (2) one or more *design goals*.

It seems plausible that Goethe’s primary design goal was long-run profit—that is, he not only wanted to maximize revenue from his manuscript, but also to find what value the publisher placed on his work. Additional design goals common in practice include:

- fairness (e.g., in school choice, Abdulkadiroğlu and Sönmez (2003), or food supply, Prendergast (2016)),
- revenue (e.g., in auctions, Myerson (1981); Binmore and Klemperer (2002); Varian (2007); Ostrovsky and Schwarz (2016)),
- simplicity/ease of use (e.g., of bidding languages, Lubin et al. (2008); Milgrom (2009, 2010); Klemperer (2010); Bichler et al. (2014), or ordinal preference elicitation, Pathak and Sönmez (2008)),
- transparency (e.g., in financial markets, Asquith et al. (2013)),
- privacy (e.g., in voting systems, McSherry and Talwar (2007)),
- speed (e.g., in algorithmic trading, Budish et al. (2015); Kyle and Lee (2017)),
- entry or competition (e.g., in auctions Levin and Smith (1994); Bajari and Hortaçsu (2003)),
- price discovery (e.g., in stock exchanges, Biais et al. (1995, 1999); Cao et al. (2000)),
- elimination of adverse selection (e.g., in ad auctions, Abraham et al. (2013); Arnosti et al. (2016), or health insurance, Finkelstein and Poterba (2004)),
- trust (e.g., in financial benchmarks, Duffie and Stein (2015)), and/or
- operation within ethical bounds (e.g., in kidney exchange, Roth et al. (2004); Roth (2007); Leider and Roth (2010)).

Marketplace designers must overcome both theoretical and practical challenges in order to achieve their goals; *complexity* characterizes the structure and difficulty of those challenges. Thus complexity of a setting is somewhat like velocity: it has both a magnitude and a direction (or, rather, a *type*). Moreover, as much as we might

want complexity to be an absolute property of design goals (a “topological property,” if you will), complexity is inextricably linked to the specific market under consideration. For instance, bidder anonymity/privacy is easy to achieve in online auctions, but is far more difficult to ensure in an auction room at Sotheby’s. Transparency and trust, on the other hand, are a headache for eBay, but are less of a problem for brick-and-mortar auction houses.⁴

3. The Auctioneer’s Parable

Milgrom’s *Discovering Prices* illustrates marketplace design in one of the most complex settings to date, showcasing how tools from economic theory, computer science, and operations research were combined in the design of Incentive Auction. Along the way, Milgrom helps us understand how to identify and address different sources of complexity.

3.1. The Setting. Radio spectrum is a limited resource: only a narrow band is useful for communication, and signals sent on overlapping channels will interfere with each other. In the US, some of the most desirable spectrum—channels on which signals can travel long distances and penetrate walls—was allocated to television broadcasters in the 1940s and 50s (Hazlett, 2008). Today, however, fewer and fewer people in the US watch broadcast television, while more and more access information through their mobile phones.

The Incentive Auction sought to buy back wireless spectrum licenses from hundreds of television broadcasters, while simultaneously selling those licenses to telecoms. The FCC hoped to clear the market as efficiently as possible, with payments from the telecoms compensating broadcasters who gave up their licenses. At the same time, the government sought to raise revenue to cover the costs associated with the process.

If all spectrum licenses were homogeneous, then the FCC could have simply run a double auction to clear the market (Wilson, 1985; McAfee, 1992). But licenses are by nature heterogeneous—they cover different areas and wavelengths (Cramton, 2002). Moreover, most broadcasters owned

local licenses, which needed to be aggregated into national networks for telecom use (Rosston, 2012; Leyton-Brown et al., 2017); this created a holdout problem. Any television broadcaster who knew that his or her license was crucial for assembling a contiguous national network could in principle have demanded compensation far above his or her true value (Kominers and Weyl, 2012; Rosston, 2012). Congress partially ameliorated holdout risk by granting the FCC the authority to move broadcasters to different channels—i.e., to *repack* them—in order to ensure that the spectrum freed up was on the same wavelengths nationwide.

But in order for the FCC to figure out which broadcasters it had to repack, it needed to know how much spectrum transfer the market would bear, and which broadcasters would sell their licenses back. Working from the rough frame of a double auction, the plan was first to run a “reverse” auction to determine which broadcast licenses could be bought back (with the remaining broadcasters being neatly repacked); then, those licenses would be auctioned off to telecoms in a “forward” (clock) auction. If the telecoms’ willingness to pay in the forward auction was not high enough to cover the spectrum acquisition cost determined in the reverse auction (plus a revenue target), the FCC would reduce the spectrum clearing target, and return to the reverse auction.

3.2. What made the Incentive Auction complex? Underlying the setting for the Incentive Auction was a computational problem: Repacking the broadcasters who did not sell back their licenses created the possibility of *interference*, whereby those broadcasters’ signals might partially or fully block each other; a given repacking would be deemed appropriate only if it created no more than “minimal” interference for every station.⁵ However, the general problem of determining when a minimal-interference repacking exists is computationally intractable—it embeds a well-known, NP-complete graph-coloring problem—making it effectively impossible to be certain which sets of broadcasters could be left

⁴But see recent work by Banksy (2018).

⁵According to the FCC criteria, a station j suffers *minimal interference* if no other station i interferes with more than 0.5% of station j ’s pre-auction audience (Leyton-Brown et al., 2017, p. 7205).

“on the air” without interference.⁶ As a consequence, there was some fundamental risk of error in identifying the efficient set of broadcast licenses to buy back.

If the efficient set of purchasable spectrum licences could not be identified with certainty, then we could at best hope to approximate the efficient allocation. But incentive-compatible mechanisms—like VCG—suggested by theory (Vickrey, 1961; Clarke, 1971; Groves, 1973; Green and Laffont, 1977; Holmström, 1979) determine prices by comparing efficient and nearly-efficient allocations; in such mechanisms, small errors in assessing efficiency get magnified into large errors in pricing (Milgrom, 2017, p. 168). The challenge was therefore to design an auction that would be (approximately) efficient and robust to imperfections in assessing repacking feasibility—yet simple enough for broadcasters to participate in. Existing economic theory provided a useful guide, but was only a starting point.

The reverse auction that Milgrom and his collaborators proposed would start by offering a high price for each broadcaster’s license. The auctioneer would then decrease the price offers in sequence. As soon as a broadcaster rejected the auctioneer’s offer, it exited the auction permanently—thus remaining on the air (Milgrom and Segal, forthcoming). To deal with concerns about interference, the auction would “freeze” the price of any broadcaster that could potentially create interference if it were to leave the auction.

The reverse auction format just described corresponds exactly to a broadcaster-proposing deferred acceptance algorithm in which broadcasters propose prices at which they would be willing to sell their licenses, and any price rejected by the auctioneer cannot be offered again (Kelso and Crawford, 1982; Hatfield and Milgrom, 2005); such a mechanism is simple to implement, and makes it particularly easy for broadcasters to figure out how to bid. Indeed, in the deferred acceptance auction, honest bidding is *obviously* optimal for unit-demand bidders (Li, 2017), in

the sense that bidders are incentivized to bid truthfully—and should be able to understand as much. As a result, under the deferred acceptance auction, bidders can figure out that truthful bidding is optimal without subtle contingent reasoning about other bidders. Additionally, the format ensures that bidders can trust the auctioneer not to tamper with prices in an undetectable way (Akbarpour and Li, 2018).

But how can we ensure that a deferred acceptance-based auction can actually clear the market? In a classic paper, Kelso and Crawford (1982) showed that the deferred acceptance algorithm can find market-clearing prices whenever goods are (gross) substitutes. If some broadcasters’ licenses are instead complements from the perspective of the FCC—in the sense that getting access to one might make others more valuable—then market-clearing prices might not even exist. One of Milgrom’s key theoretical contributions in *Discovering Prices* (pp. 190–194) is that if goods are, in a formal sense “near-substitutes,” then the deferred acceptance algorithm will reach prices that are nearly efficient, and the resulting allocation will be close to efficient, as well.

To be a bit more precise: Milgrom considers the structure of the FCC’s preferences over broadcast licenses. The FCC wanted to buy back licenses in a way that would maximize the value of the broadcasters who kept their licenses (and thus remained on the air) in the end. That is, the set of *rejected broadcasters* that the FCC does not acquire licenses from should be those with highest value for operating as stations.⁷ From an auction design perspective, it would be ideal if the FCC’s preferences were (*one-for-one*) *substitutable*, meaning that if, at a given price profile, the FCC chooses to leave set of broadcasters S on the air but some station $s \in S$ indicates that it will be willing to sell its license at a new, lower price, then the FCC will still want to leave the broadcasters in $S \setminus \{s\}$ on the air.

Unfortunately, because of the practical constraints on repacking, the FCC’s preferences were

⁶A problem is *in (complexity class) NP* if it can be solved in nondeterministic polynomial time; a problem is *NP-complete* if it is in NP and all problems in NP can be reduced to it (Sipser, 2012). At least assuming a mathematical conjecture widely believed to be true (Cook, 2006), solutions to problems in NP cannot, in general, be found efficiently (although they can be verified quickly).

⁷For technical reasons, it is typically easier to work with the set of broadcasters whose offers to sell back their licenses are rejected, rather than the set whose are accepted.

not substitutable. It was possible, for example, that when a small-bandwidth broadcaster reduced its buyout price, the FCC would want to reject a large-bandwidth broadcaster and instead buy back licenses from a set of currently rejected small-bandwidth stations. To get around this problem, Milgrom showed how to approximate non-substitutable preferences, like those of the FCC, with substitutable ones. The approximation gives rise to an intuitive and robust “index of substitutability” ρ that is equal to 1 whenever preferences are exactly substitutable, and approaches 0 when stations become perfectly complementary.

Milgrom shows that the worst-case efficiency performance of the deferred acceptance auction (across all valuations) is precisely ρ (Proposition 4.6 of Milgrom (2017); see also de Vries and Vohra (2019)).⁸ That is, the more substitutable the FCC’s preferences are according to Milgrom’s index, the closer the worst-case performance of the deferred acceptance auction is to the fully efficient outcome. By making repacking rules quite flexible, the FCC sought to make licenses as substitutable as possible (Kominers and Weyl, 2012; Rosston, 2012). In simulations, the deferred acceptance auction achieved around 95% efficiency (Leyton-Brown et al., 2017).

Even with deferred acceptance auctions assured to produce nearly efficient outcomes, the repacking problem would still need to be solved at every round of the auction clock. Thus, the design team developed state-of-the-art infrastructure for finding feasible repackings quickly, based on newly-developed *deep optimization* techniques (Newman et al., 2017). The new computational infrastructure ensured that the feasibility-checking software did not timeout too frequently—and it would never make Type I (“false positive”) errors.⁹

While the reverse auction was a “key and innovative part of the FCC’s Incentive Auction” (Cramton et al., 2015, p. 1), the forward clock

auction was also quite intricate,¹⁰ as were the auction-closing rule and the “assignment round” in which different frequencies were allocated to telecoms.

And with so many important details to get right, it would be unlikely for everything to go entirely according to plan. *Discovering Prices* went to press as the assignment round was still running; that phase used a Vickrey auction to set final prices for particular bandwidths. It is well-known that when bidders have complementary preferences over frequencies, some Vickrey auction equilibria could result in low revenues (Milgrom, 2004; Ausubel and Milgrom, 2006; Marszalec, 2018). And indeed, this is exactly what appears to have happened: Ausubel et al. (2017) estimate that successful coordination by the bidders cost the assignment round around \$5 million in revenue (admittedly just 3.5% of the round’s revenue, and less than a tenth of a percent of the overall auction revenue). Meanwhile, some private equity funds appear to have coordinated ownership of “Class A” broadcast licenses in the reverse auction, reducing supply and substantially raising the FCC’s acquisition costs (Doraszelski et al., 2017).

4. Three Types of Complexity

The story of the Incentive Auction as told in *Discovering Prices* vividly illustrates at least three types of complexity that also arise in other design settings. First, how should marketplace designers specify and trade-off various design goals? Second, how can a marketplace designer overcome difficulties of eliciting information from and/or providing information to market participants? Third—perhaps least familiar to economists—how should marketplace designer tackle fundamental computational constraints? As we already saw in Section 3, computational complexity at the heart of spectrum reallocation was a large part of what made the setting so challenging.

⁸The result uses insights from the theory of *matroids* (see also Korte and Hausmann (1978); Ostrovsky and Paes Leme (2015)).

⁹That is, the feasibility-checker would only deem a set of broadcasters non-interfering if a permissible repacking could be found. But by nature, there was still some possibility of Type II errors (“false negatives”): for some configurations of broadcasters, it might not be possible to find a feasible repacking sufficiently quickly; in that case, the system would declare the set of broadcasters potentially interfering.

¹⁰It featured *intra-round bidding* to avoid bidder exposure, as well as *condition reserves* to promote competition.

4.1. Specification complexity. In many market design settings, it can be difficult to translate abstract design goals into practical objectives. And even once we have done so, our design goals may be hard—or even impossible—to achieve simultaneously. When goals conflict and/or are difficult to express, simply specifying *what* to prioritize becomes a design challenge all of its own; *specification complexity* reflects these decisions and tradeoffs.

Specification complexity in spectrum reallocation. The FCC’s early remarks on its goals for repurposing spectrum were extremely abstract. For example, the 2010 National Broadband Plan recommended that the policymakers “ensure greater transparency concerning spectrum allocation and utilization” and “expand incentives and mechanisms to reallocate or repurpose spectrum” (FCC, 2010, p. 75). One can imagine that there are many objective functions that would satisfy such an abstract specification. The exercise of making the design goals more precise was the first principal complexity in the design of the Incentive Auction. Following a wide-ranging consultation, the policymakers decided that the goals of the auction would be: (1) to buy out hundreds of TV broadcasters with heterogeneous licenses, (2) to reallocate their licenses to telecoms as efficiently as possible, and (3) to ensure a reasonable amount of revenue for the government. But even these goals were not precise enough to serve as guidance for how the auction should be run. If the market were one-sided (e.g., if the FCC only had to sell available spectrum), then the problem would have been easier—the FCC could have designed an efficient auction and simply increased the amount of spectrum on sale until the revenue target was hit. But designing efficient centralized two-sided markets with privately informed participants typically requires a subsidy (Myerson and Satterthwaite, 1983). Thus, the revenue vs. efficiency trade-off is much starker in broadcast spectrum reallocation settings than in typical spectrum auctions (Loertscher et al., 2015).

The auctioneers resolved the objective specification complexity in two ways. First, in order to increase competition between broadcasters, spectrum was homogenized by allowing repacking. Second, the need to figure out how much spectrum to clear while ensuring that the revenue

target was hit led the designers to run the forward and reverse auctions sequentially, adjusting the clearing target as they went along.

Specification complexity in other settings. Even settings with few design goals often force the market designer to make difficult decisions about ill-specified objectives. Take efficiency, for example. We often think of “efficient outcomes” as being those that maximize allocative surplus (the “Kaldor–Hicks criterion”). In Goethe’s case, we might think of efficiency as purely allocative: the goal was simply to deliver the manuscript to the publisher who valued it the most. But efficiency really means maximizing welfare—and thus our interpretation of efficiency as a design goal depends on the welfare function. If we have distributional preferences, then efficiency might involve distorting the allocation—or even wasting resources—in order to improve the outcomes of poorer people in the market (Rawls (1971); Weitzman (1977); see also Dworzak & Kominers & Akbarpour (2018)). Likewise, design goals such as “fairness” and “transparency” depend crucially on the market in question, and may be complex to define in practice.

Even once we understand our design goals individually, balancing among them can be complex. Sometimes, trade-offs arise from formal impossibility results: Voting rules satisfying fairly mild desiderata end up being dictatorial (Arrow, 1951; Gibbard, 1973; Satterthwaite, 1975). In school choice and similar matching settings, it is typically impossible to simultaneously ensure (Pareto) efficiency, strategy-proofness, and elimination of justified envy (Abdulkadiroğlu and Sönmez, 2003; Abdulkadiroğlu et al., 2009; Kesten, 2010). Other times, some of our design goals pin down a unique rule or a class of rules—but all the mechanisms in that class might fail to satisfy some other design goal. For example, in an object allocation setting with independent private values and transfers, if we always want to choose efficient outcomes in a strategy-proof way, we must use a Groves mechanism (Groves, 1973; Green and Laffont, 1977; Holmström, 1979; Williams, 1999). But Groves mechanisms are not budget-balanced and might be susceptible to collusion (Klemperer, 2004; Ausubel and Milgrom, 2006). Another common example arises when

governments run spectrum auctions and set efficiency and revenue as their two key goals—it is well-known that these objectives conflict even in single-object auctions (Myerson, 1981). Yet more complex trade-offs arise from practical considerations: in ecosystem services markets in developing countries, for example, market designers must balance development targets with environmental concerns (Jayachandran et al., 2017; Salzman et al., 2018; Wunder et al., 2018; Teytelboym, 2019).

Specification complexity can also arise in the seemingly simple process of choosing which transactions should be allowed. When goods are commodities, every transaction is like every other (for a given transaction volume). But in markets with substantial heterogeneity, it may be hard to determine which types of transactions to “conflate” (i.e., treat as identical), and which to keep distinct. These sorts of questions arise in ride-sharing (how local should pricing be?), online advertising (how fine-grained should audience targeting be?), and most spectrum auctions (which bands should be sold together?). By conflating sets of slightly different transactions, the market designer might be able to increase market thickness, reduce cherry-picking, and make the marketplace safer for participants (Levin and Milgrom, 2010; Milgrom, 2010). But at the same time, conflation might harm efficiency or other allocative goals by limiting the marketplace’s ability to respond to heterogeneity in preferences.

4.2. Informational complexity. Even after the designer’s objectives are clear, information frictions often stand in the way. The market needs to somehow “discover” participants’ preferences. And before that, participants need to acquire and process information in order to access the market and form their preferences. *Informational complexity* reflects the difficulties that arise in aggregating information from market participants and/or in providing information to them.

Informational complexity in spectrum reallocation. Figuring out how much to compensate broadcasters for selling their licenses was a daunting challenge. If repacking had not been possible,

then broadcaster holdout on its own would have made reassembly impossible: there was likely so much complementarity between TV stations that the FCC would have had no chance of determining the appropriate compensation to take a nationwide band of them off-air (Mailath and Postlewaite, 1990). Fortunately, the ability to repack TV stations created enough competition to be able to think of the market as approximately competitive (Cramton et al., 2015).¹¹ While the forward auction rules could, in principle, build on previous package auctions for spectrum, running the reverse auction to clear the TV stations off-air was no trivial matter. The natural strategy-proof way to find the efficient outcome—running a VCG auction—would have been confusing for the broadcasters and, in the presence of possible complementarities, would have had undesirable properties (Rothkopf, 2007). Informational complexity therefore remained in two forms: how to elicit broadcasters’ valuations as accurately as possible and how to give broadcasters clear and credible signals about how to bid. The deferred acceptance auction elegantly tackled this informational complexity; meanwhile, Milgrom’s results on the robustness of the deferred acceptance auction in the presence of complementarities gave some reassurance about how closely the outcome might approximate efficiency.

Informational complexity in other settings. In many settings, of course, market-clearing prices are enough to both aggregate preferences and guide market-clearing. But communication frictions, preference heterogeneity, and/or complementarities among goods may make it hard to determine market prices (Hurwicz, 1960; Mount and Reiter, 1974; Nisan and Segal, 2006; Segal, 2007). Pricing is also difficult whenever utility is not fully transferable, either because of societal preferences or incomplete markets (Fleiner et al., forthcoming). Additionally, efficient pricing may conflict with other design goals.¹²

When prices are unavailable to aid decentralized market-clearing, we sometimes develop

¹¹However, as we noted briefly in Section 3.2, Doraszelski et al. (2017) estimate that strategic supply reduction did impact the market, and may have increased FCC’s acquisition costs by 22%.

¹²For example, in auctions to sell complements, Vickrey prices sit outside the core, result in low revenue, and encourage collusion. In response, many different core-pricing rules have been proposed (Day and Milgrom, 2008; Erdil and Klemperer, 2010; Day and Cramton, 2012; Bünz et al., 2018).

centralized mechanisms in which we elicit participants’ preferences directly. But here again complexity can arise: Many systems for assigning children to schools, for example, ask parents to rank-order the schools their children may attend. If the assignment mechanism is strategy-proof, then it is straightforward for parents to submit preferences—they simply list the schools they find acceptable in true preference order. But submitting even a simple list may be difficult for parents if the assignment system is manipulable; moreover, in that case, the designer struggles to determine how the submitted rank-order lists relate to parents’ true preferences. Moreover, even if a marketplace can elicit true preferences from individual agents, it may be unable to prevent more complicated manipulations. For example, Vickrey auctions can not only be manipulated by coordinated misreporting (i.e., they are not “collusion-proof”), but also by using shill bidders (i.e., they are not “false-name-proof”; see [Yokoo et al. \(2004\)](#)).

Complexity might also arise when the designer needs to elicit combinatorial preferences over a large set of objects. In contexts like display ad auctions ([Lahaie et al., 2008](#)), electricity markets ([Schnizler and Neumann, 2007](#)), and business school course allocation ([Sönmez and Ünver, 2010](#); [Budish, 2011](#); [Budish and Cantillon, 2012](#)), some goods might be complements, while others are substitutes—meaning that rank-order lists are insufficient to characterize participants’ preferences. Yet submitting high-dimensional combinatorial preferences is infeasible, necessitating the design of simpler bidding languages ([Milgrom, 2009](#); [Bichler et al., 2014](#); [Budish and Kessler, 2016](#); [Budish et al., 2016](#)).

In still other markets, complexity can arise from the need to provide information to help participants learn their own preferences. It can be nontrivial, for example, to enable parents to understand the trade-offs between different schooling options for their children ([Hastings et al., 2007](#); [Corcoran et al., 2018](#)). And even structured marketplace exchanges may lack ratings or certification systems that would help participants

identify high-quality trading partners ([Pallais, 2014](#); [Li et al., 2016](#)).¹³

4.3. Computational complexity. Questions of how to specify objectives and structure information have been central in market(place) design—and economics, more broadly—for a long time. But many market design settings have a further layer: *computational complexity*, under which fundamental computational constraints impede the designer’s ability to achieve his or her goals.¹⁴ Computational problems affect the extent of specification and informational complexity, and can also introduce independent practical problems.

Computational complexity in spectrum reallocation. As we have already discussed, the problem of determining which sets of broadcasters could be left on the air with minimal interference was computationally intractable. This computational complexity constrained the design of the Incentive Auction substantially—in particular, it made computing efficient allocations impractical, which ruled out VCG and similar mechanisms. In VCG, errors in solving the repacking problem are amplified by interdependence of prices and by the need to compute prices in one step. By adjusting prices gradually, the deferred acceptance auction was less sensitive to imperfections in feasibility checking and more likely to find market-clearing prices (and thus an efficient allocation). Even so, advances in deep optimization were required in order to ensure that the repacking problems could be solved quickly and accurately enough that the auctioneer would know which prices could be reduced at each stage.

Computational complexity in other settings. More generally, we know from computer science that some types of computation (say, the “set cover” problem or finding (Bayesian) Nash equilibria) are difficult to execute quickly. On the other hand, it is straightforward to run English auctions to find efficient allocations of heterogeneous, substitutable items ([Kelso and Crawford, 1982](#); [Demange et al., 1986](#); [Gul and Stacchetti,](#)

¹³For further discussion, see Sections 4 and 5 of [Kominers \(2018\)](#).

¹⁴Our usage of the term “computational complexity” is not too distant from its use in theoretical computer science. Complexity theory characterizes the inherent difficulty of different computational problems; our broader usage here also includes practical computational challenges.

2000). But when markets are dynamic or feature substantial heterogeneity, computing market-clearing prices is often very difficult (Nisan and Segal, 2006; Gonczarowski et al., 2014). Likewise, in two-sided matching markets we can find stable outcomes via the deferred acceptance algorithm (Gale and Shapley, 1962). But in markets with combinatorial preferences, it can be computationally intractable to compute a stable outcome or even check whether one exists (McDermid and Manlove, 2010; Manlove, 2013).

Sometimes, we can “relax” computationally complex problems to solve them faster and/or approximately (see, e.g., Nguyen and Vohra (2018, forthcoming); Nguyen et al. (2019))—yet practical concerns may rule such relaxations out. For example, computing fully efficient allocations in electricity markets is computationally difficult, yet if we seek approximate solutions we might end up with mismatches in supply and demand that cause blackouts (Cramton, 2017).

5. Outlook

So what have we discovered from *Prices*? Modern marketplace design increasingly wrestles with complexity; as it does so, we need novel, tailor-made theory as well as supporting infrastructure. Complexity has real economic meaning—and can take multiple forms.

Discovering Prices teaches us how to map and isolate different forms of complexity—to pin down exactly where the standard economic toolkit can be applied, where existing methods can be extended, and where entirely new tools are required. In doing so, Milgrom (2017) proves that marketplace design can work in increasingly complex settings. Briefly returning to the ride-sharing example from the introduction. Ride-share marketplaces face specification complexity in determining how to trade-off efficiency and revenue goals. They must then elicit information about riders’ and drivers’ multidimensional preferences over cost, direction, and speed. Computational complexity plays a role, too, as optimal allocations and travel schedules must be computed quickly, in real time. Once isolated, these different forms of complexity become approachable both in theory in and practice (see, for example, Manna

and Prestwich (2014); Masoud and Jayakrishnan (2017); Ma et al. (2018); Rheingans-Yoo et al. (2019)). And surely even more complex settings will yield to our design efforts in years to come.

Marketplace design in increasingly complex settings may bring political and societal challenges. As Hitzig (2018) argues, the Incentive Auction is a signal that a future market designer advising the government might no longer be able to act simply as competent technocrat or a humble “economic engineer” (Roth, 2002), but would have to become “a craftsman [and] a technologist who implements as well as designs, and who creates as well as conceives”. As market design moves into complex settings, it becomes less clear how designs can be audited and how designers can be scrutinized. Market designers must not lose sight of the need to maintain public trust. But checks on the remit of market designers should not necessarily place limits on the expanding role of market design. At a time when politicians say that “the people . . . have had enough of experts,”¹⁵ Milgrom (2017) gives us hope that market designers will still be called upon to showcase their craft.

References

- Abdulkadiroğlu, Atila and Tayfun Sönmez, “School choice: A mechanism design approach,” *American economic review*, 2003, 93 (3), 729–747.
- Abdulkadiroğlu, Atila and Tayfun Sönmez, “School Choice: A Mechanism Design Approach,” *American Economic Review*, 2003, 93, 729–747.
- , Parag A. Pathak, and Alvin E. Roth, “Strategyproofness versus Efficiency in Matching with Indifferences: Redesigning the NYC High School Match,” *American Economic Review*, 2009, 99, 1954–1978.
- Abraham, Ittai, Susan Athey, Moshe Babaioff, and Michael Grubb, “Peaches, lemons, and cookies: Designing auction markets with dispersed information,” in “Proceedings of the Fourteenth ACM Conference on Electronic Commerce” 2013, pp. 7–8.

¹⁵See, for example, Mance (2016).

- Aggarwal, Gagan, Ashish Goel, and Rajeev Motwani**, “Truthful auctions for pricing search keywords,” in “Proceedings of the 7th ACM conference on Electronic Commerce” ACM 2006, pp. 1–7.
- Akbarpour, Mohammad and Shengwu Li**, “Credible mechanisms,” in “Proceedings of the 2018 ACM Conference on Economics and Computation” ACM 2018, pp. 371–371.
- Arnosti, Nick, Marissa Beck, and Paul Milgrom**, “Adverse selection and auction design for Internet display advertising,” *American Economic Review*, 2016, 106 (10), 2852–66.
- Arrow, Kenneth J.**, *Social choice and individual values*, John Wiley & Sons, 1951.
- and **Leonid Hurwicz**, “On the stability of the competitive equilibrium, I,” *Econometrica: Journal of the Econometric Society*, 1958, pp. 522–552.
- Asquith, Paul, Thom Covert, and Parag A. Pathak**, “The effects of mandatory transparency in financial market design: Evidence from the corporate bond market,” 2013. NBER Working Paper No. 19417.
- Athey, Susan and Denis Nekipelov**, “A structural model of sponsored search advertising auctions,” in “Sixth Ad Auctions Workshop,” Vol. 15 2010.
- and **Glenn Ellison**, “Position Auctions with Consumer Search,” *Quarterly Journal of Economics*, 2011, 126, 1213–1270.
- Ausubel, Lawrence M. and Paul R. Milgrom**, “The lovely but lonely Vickrey auction,” in Peter Cramton, Yoav Shoham, and Richard Steinberg, eds., *Combinatorial Auctions*, MIT Press, 2006, pp. 17–40.
- Ausubel, Lawrence M, Christina Aperjis, and Oleg Baranov**, “Market design and the FCC incentive auction,” 2017. Working Paper.
- Bajari, Patrick and Ali Hortacısu**, “The winner’s curse, reserve prices, and endogenous entry: Empirical insights from eBay auctions,” *RAND Journal of Economics*, 2003, 34 (2), 329–355.
- Banksy**, “Love is in the Bin,” 2018. Originally auctioned as “Canvas of Girl With Balloon”.
- Biais, Bruno, Pierre Hillion, and Chester Spatt**, “An empirical analysis of the limit order book and the order flow in the Paris Bourse,” *Journal of Finance*, 1995, 50 (5), 1655–1689.
- , ———, and ———, “Price discovery and learning during the preopening period in the Paris Bourse,” *Journal of Political Economy*, 1999, 107 (6), 1218–1248.
- Bichler, Martin, Jacob Goeree, Stefan Mayer, and Pasha Shabalin**, “Spectrum auction design: Simple auctions for complex sales,” *Telecommunications Policy*, 2014, 38 (7), 613–622.
- Binmore, Ken and Paul Klemperer**, “The biggest auction ever: the sale of the British 3G telecom licences,” *The Economic Journal*, 2002, 112 (478), C74–C96.
- Borges, Jorge Luis**, “Lottery in Babylon,” *Prairie Schooner*, 1959, 33 (3), 203–207. Translated by J. M. Fein.
- Börner, Lars and John William Hatfield**, “The Design of Debt-Clearing Markets: Clearinghouse Mechanisms in Pre-Industrial Europe,” Mimeo November 2012.
- and ———, “The Design of Debt-Clearing Markets: Clearinghouse Mechanisms in Preindustrial Europe,” *Journal of Political Economy*, 2017, 125 (6), 1991–2037.
- Budish, Eric**, “The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes,” *Journal of Political Economy*, 2011, 119 (6), 1061–1103.
- and **Estelle Cantillon**, “The multi-unit assignment problem: Theory and evidence from course allocation at Harvard,” *American Economic Review*, 2012, 102 (5), 2237–2271.
- and **Judd B. Kessler**, “Bringing Real Market Participants’ Real Preferences into the Lab: An Experiment that Changed the Course Allocation Mechanism at Wharton,” 2016. NBER Working Paper No. 22448.
- , **Gérard P Cachon, Judd B Kessler, and Abraham Othman**, “Course match: A large-scale implementation of approximate competitive equilibrium from equal incomes for combinatorial allocation,” *Operations Research*, 2016, 65 (2), 314–336.
- , **Peter Cramton, and John Shim**, “The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response,” *Quarterly Journal of Economics*, 2015, 130 (4), 1547–1621.
- Bünz, Benedikt, Benjamin Lubin, and**

- Sven Seuken**, “Designing core-selecting payment rules: A computational search approach,” 2018.
- Cao, Charles, Eric Ghysels, and Frank Hatheway**, “Price discovery without trading: Evidence from the Nasdaq preopening,” *Journal of Finance*, 2000, 55 (3), 1339–1365.
- Clarke, Edward H.**, “Multipart Pricing of Public Goods,” *Public Choice*, 1971, 11 (1), 17–33.
- Cook, Stephen**, “The P versus NP problem,” in James Carlson, Arthur Jaffe, and Andrew Wiles, eds., *The Millennium Prize Problems*, Clay Mathematics Institute, 2006, pp. 87–104.
- Corcoran, Sean P., Jennifer L. Jennings, Sarah R. Cohodes, and Carolyn Sattin-Bajaj**, “Leveling the Playing Field for High School Choice: Results from a Field Experiment of Informational Interventions,” 2018. NBER Working Paper No. 24471.
- Cramton, Peter**, “Spectrum auctions,” in Martin Cave, Sumit Majumdar, and Ingo Vogelsang, eds., *Handbook of Telecommunications Economics*, Elsevier, 2002, pp. 605–639.
- , “Electricity market design,” *Oxford Review of Economic Policy*, 2017, 33 (4), 589–612.
- , **Hector Lopez, David Malec, and Pacharasut Sujarittanonta**, “Design of the Reverse Auction in the FCC Incentive Auction,” 2015. Expert Report in response to Comment Public Notice FCC 14-191.
- Dantzig, George B.**, “Discrete-variable extremum problems,” *Operations research*, 1957, 5 (2), 266–288.
- Day, Robert W. and Paul Milgrom**, “Core-selecting package auctions,” *International Journal of Game Theory*, 2008, 36 (3), 393–407.
- **and Peter Cramton**, “Quadratic core-selecting payment rules for combinatorial auctions,” *Operations Research*, 2012, 60 (3), 588–603.
- de Vries, Sven and Rakesh V. Vohra**, “Matroidal Approximations of Independence Systems,” arXiv:1906.06217 June 2019.
- Demange, Gabrielle, David Gale, and Marilda Sotomayor**, “Multi-item auctions,” *Journal of Political Economy*, 1986, 94 (4), 863–872.
- Doraszelski, Ulrich, Katja Seim, Michael Sinkinson, and Peichun Wang**, “Ownership concentration and strategic supply reduction,” 2017. NBER Working Paper No. 23034.
- Duffie, Darrell and Jeremy C. Stein**, “Reforming LIBOR and other financial market benchmarks,” *Journal of Economic Perspectives*, 2015, 29 (2), 191–212.
- Dworczak, P. & S. D. Kominers & M. Akbarpour**, “Redistribution through Markets,” 2018. Stanford Graduate School of Business Working Paper.
- Edelman, Benjamin, Michael Ostrovsky, and Michael Schwarz**, “Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords,” *American Economic Review*, 2007, 97, 242–259.
- Eisenmann, Thomas and Scott Duke Kominers**, “Making Markets,” 2018. Harvard Business School Technical Note 818-096.
- Erdil, Aytek and Paul Klemperer**, “A new payment rule for core-selecting package auctions,” *Journal of the European Economic Association*, 2010, 8 (2-3), 537–547.
- FCC**, *Connecting America: The National Broadband Plan*, Federal Communication Commission, 2010.
- Finkelstein, Amy and James Poterba**, “Adverse selection in insurance markets: Policyholder evidence from the UK annuity market,” *Journal of Political Economy*, 2004, 112 (1), 183–208.
- Fleiner, Tamás, Ravi Jagadeesan, Zsuzsanna Jankó, and Alexander Teytelboym**, “Trading Networks with Frictions,” *Econometrica*, forthcoming.
- Gale, David and Lloyd S. Shapley**, “College admissions and the stability of marriage,” *American Mathematical Monthly*, 1962, 69, 9–15.
- Gibbard, Allan**, “Manipulation of Voting Schemes: A General Result,” *Econometrica*, 1973, 41 (4), 587–601.
- Gonczarowski, Yannai A., Noam Nisan, Rafail Ostrovsky, and Will Rosenbaum**, “A stable marriage requires communication,” in “Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms” SIAM 2014, pp. 1003–1017.
- Green, Jerry and Jean-Jacques Laffont**, “Characterization of Satisfactory Mechanisms

- for the Revelation of Preferences for Public Goods,” *Econometrica*, 1977, 45, 427–438.
- Groves, Theodore**, “Incentives in Teams,” *Econometrica*, 1973, 41 (4), 617–631.
- Gul, Faruk and Ennio Stacchetti**, “Walrasian equilibrium with gross substitutes,” *Journal of Economic theory*, 1999, 87 (1), 95–124.
- and ———, “The English auction with differentiated commodities,” *Journal of Economic theory*, 2000, 92 (1), 66–95.
- Hastings, Justine S., Richard Van Weelden, and Jeffrey Weinstein**, “Preferences, information, and parental choice behavior in public school choice,” 2007. NBER Working Paper No. 12995.
- Hatfield, John William and Paul R. Milgrom**, “Matching with Contracts,” *American Economic Review*, 2005, 95, 913–935.
- , **Scott Duke Kominers, Alexandru Nichifor, Michael Ostrovsky, and Alexander Westkamp**, “Stability and competitive equilibrium in trading networks,” *Journal of Political Economy*, 2013, 121, 966–1005.
- Hazlett, Thomas W.**, “Optimal abolition of FCC spectrum allocation,” *Journal of Economic Perspectives*, 2008, 22 (1), 103–128.
- Hitzig, Zoë**, “Economist as craftsperson in the FCC’s “Incentive Auction”,” September 2018. HISRECO Conference.
- Holmström, Bengt**, “Groves’ Scheme on Restricted Domains,” *Econometrica*, 1979, 47 (5), 1137–1144.
- Hurwicz, L.**, “Optimality and Informational Efficiency in Resource Allocation Processes,” in K.J. Arrow, S. Karlin, and P. Suppes, eds., *Mathematical Methods in the Social Sciences*, Stanford University Press, 1960, pp. 27–46.
- Jayachandran, Seema, Joost de Laat, Eric F. Lambin, Charlotte Y. Stanton, Robin Audy, and Nancy E. Thomas**, “Cash for carbon: A randomized trial of payments for ecosystem services to reduce deforestation,” *Science*, 2017, 357 (6348), 267–273.
- Kelso, Alexander S. and Vincent P. Crawford**, “Job Matching, Coalition Formation, and Gross Substitutes,” *Econometrica*, 1982, 50, 1483–1504.
- Kesten, Onur**, “School choice with consent,” *Quarterly Journal of Economics*, 2010, 125, 1297–1348.
- Klemperer, Paul**, *Auctions: Theory and Practice*, Princeton University Press, 2004.
- , “The product-mix auction: A new auction design for differentiated goods,” *Journal of the European Economic Association*, 2010, 8 (2-3), 526–536.
- Kominers, Scott Duke**, “Good Markets (Really Do) Make Good Neighbors,” *SIGecom Exchanges*, 2018, 16 (2), 12–26.
- , **Alexander Teytelboym, and Vincent P. Crawford**, “An invitation to market design,” *Oxford Review of Economic Policy*, 2017, 33 (4), 541–571.
- and **E. Glen Weyl**, “Holdout in the Assembly of Complements: A Problem for Market Design,” *American Economic Review*, 2012, 102, 360–365.
- Korte, Bernhard and Dirk Hausmann**, “An analysis of the greedy heuristic for independence systems,” in B. Alspach, P. Hell, and D.J. Miller, eds., *Annals of Discrete Mathematics*, Vol. 2 of *Algorithmic Aspects of Combinatorics*, Elsevier, 1978, pp. 65–74.
- Kyle, Albert S. and Jeongmin Lee**, “Toward a fully continuous exchange,” *Oxford Review of Economic Policy*, 2017, 33 (4), 650–675.
- Lahaie, Sébastien, David C. Parkes, and David M. Pennock**, “An Expressive Auction Design for Online Display Advertising,” in “AAAI” 2008, pp. 108–113.
- Leider, Stephen and Alvin E. Roth**, “Kidneys for sale: Who disapproves, and why?,” *American Journal of Transplantation*, 2010, 10, 1221–1227.
- Levin, Dan and James L Smith**, “Equilibrium in auctions with entry,” *The American Economic Review*, 1994, 84 (3), 585–599.
- Levin, Jonathan and Paul Milgrom**, “Online advertising: Heterogeneity and conflation in market design,” *American Economic Review*, 2010, 100 (2), 603–07.
- Leyton-Brown, Kevin, Paul Milgrom, and Ilya Segal**, “Economics and computer science of a radio spectrum reallocation,” *Proceedings of the National Academy of Sciences*, 2017, 114 (28), 7202–7209.
- Li, Lingfang Ivy, Steven Tadelis, and Xiaolan Zhou**, “Buying reputation as a signal of quality: Evidence from an online marketplace,”

2016. NBER Working Paper No. 22584.
- Li, Shengwu**, “Obviously strategy-proof mechanisms,” *American Economic Review*, 2017, *107* (11), 3257–87.
- Loertscher, Simon, Leslie M. Marx, and Tom Wilkening**, “A Long Way Coming: Designing Centralized Markets with Privately Informed Buyers and Sellers,” *Journal of Economic Literature*, 2015, *53* (4), 857–97.
- Lubin, Benjamin, Adam I. Juda, Ruggiero Cavallo, Sébastien Lahaie, Jeffrey Shneidman, and David C. Parkes**, “ICE: An expressive iterative combinatorial exchange,” *Journal of Artificial Intelligence Research*, 2008, *33*, 33–77.
- Ma, Hongyao, Fei Fang, and David C. Parkes**, “Spatio-Temporal Pricing for Ridesharing Platforms,” 2018. Working Paper.
- Mailath, George J. and Andrew Postlewaite**, “Asymmetric information bargaining problems with many agents,” *Review of Economic Studies*, 1990, *57* (3), 351–367.
- Mance, Henry**, “Britain has had enough of experts, says Gove,” *Financial Times* June 2016.
- Manlove, David F.**, *Algorithmics of Matching Under Preferences*, Vol. 2 of *Series on Theoretical Computer Science*, World Scientific, 2013.
- Manna, Carlo and Steve Prestwich**, “Online stochastic planning for taxi and ridesharing,” in “2014 IEEE 26th International Conference on Tools with Artificial Intelligence (ICTAI) IEEE 2014, pp. 906–913.
- Marszalec, Daniel**, “Fear not the simplicity—An experimental analysis of auctions for complements,” *Journal of Economic Behavior & Organization*, 2018, *152*, 81–97.
- Masoud, Neda and R. Jayakrishnan**, “A real-time algorithm to solve the peer-to-peer ride-matching problem in a flexible ridesharing system,” *Transportation Research Part B: Methodological*, 2017, *106*, 218–236.
- McAfee, R. Preston**, “A dominant strategy double auction,” *Journal of Economic Theory*, 1992, *56* (2), 434–450.
- McDermid, Eric J. and David F. Manlove**, “Keeping partners together: algorithmic results for the hospitals/residents problem with couples,” *Journal of Combinatorial Optimization*, 2010, *19* (3), 279–303.
- McSherry, Frank and Kunal Talwar**, “Mechanism design via differential privacy,” in “Foundations of Computer Science, 2007. FOCS’07. 48th Annual IEEE Symposium on” IEEE 2007, pp. 94–103.
- Milgrom, Paul**, *Putting Auction Theory to Work*, Cambridge University Press, 2004.
- and **Bruno Strulovici**, “Substitute goods, auctions, and equilibrium,” *Journal of Economic theory*, 2009, *144* (1), 212–247.
- Milgrom, Paul R.**, “Assignment Messages and Exchanges,” *American Economic Journal: Microeconomics*, 2009, *1*, 95–113.
- , “Simplified Mechanisms with an Application to Sponsored-Search Auctions,” *Games and Economic Behavior*, 2010, *70*, 62–70.
- , *Discovering Prices: Auction Design in Markets with Complex Constraints*, Columbia University Press, 2017.
- and **Ilya Segal**, “Clock Auctions and Radio Spectrum Reallocation,” *Journal of Political Economy*, forthcoming.
- Moldovanu, Benny and Manfred Tietzel**, “Goethe’s second-price auction,” *Journal of Political Economy*, 1998, *106* (4), 854–859.
- Mount, Kenneth and Stanley Reiter**, “The informational size of message spaces,” *Journal of Economic Theory*, June 1974, *8* (2), 161–192.
- Myerson, Roger B.**, “Optimal auction design,” *Mathematics of Operations Research*, 1981, *6*, 58–73.
- and **Mark A. Satterthwaite**, “Efficient mechanisms for bilateral trading,” *Journal of Economic Theory*, 1983, *29*, 265–281.
- Newman, Neil, Alexandre Fréchet, and Kevin Leyton-Brown**, “Deep optimization for spectrum repacking,” *Communications of the ACM*, 2017, *61* (1), 97–104.
- Nguyen, Thành and Rakesh Vohra**, “Near-Feasible Stable Matchings with Couples,” *American Economic Review*, 2018.
- and ———, “Stable Matching with Proportionality Constraints,” *Operations Research*, forthcoming.
- , **Hai Nguyen, and Alexander Teytelboym**, “Stability in matching markets with complex constraints,” in “Proceedings of the 2019 ACM Conference on Economics and Computation” ACM 2019, pp. 61–61.

- Nisan, Noam and Ilya Segal**, “The communication requirements of efficient allocations and supporting prices,” *Journal of Economic Theory*, 2006, 129 (1), 192–224.
- Ostrovsky, Michael and Michael Schwarz**, “Reserve Prices in Internet Advertising Auctions: A Field Experiment,” 2016. Stanford Graduate School of Business Working Paper.
- and **Renato Paes Leme**, “Gross substitutes and endowed assignment valuations,” *Theoretical Economics*, 2015, 10 (3), 853–865.
- Pallais, Amanda**, “Inefficient Hiring in Entry-Level Labor Markets,” *American Economic Review*, 2014, 104 (11), 3565–3599.
- Pathak, Parag A. and Tayfun Sönmez**, “Leveling the Playing Field: Sincere and Sophisticated Players in the Boston Mechanism,” *American Economic Review*, 2008, 98, 1636–1652.
- Prendergast, Canice**, “The Allocation of Food to Food Banks,” Technical Report, University of Chicago 2016.
- Rawls, John**, *A Theory of Justice*, Belknap, Harvard University Press, 1971.
- Rheingans-Yoo, Duncan, Scott Duke Kominers, Hongyao Ma, and David C. Parkes**, “Ridesharing with Driver Location Preferences,” in “Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19” International Joint Conferences on Artificial Intelligence Organization 7 2019, pp. 557–564.
- Rosston, Gregory**, “Incentive Auctions,” *Communications of the ACM*, 2012, 55, 24–26.
- Roth, Alvin E.**, “The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory,” *Journal of Political Economy*, 1984, 92, 991–1016.
- , “The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics,” *Econometrica*, 2002, 70, 1341–1378.
- , “Repugnance as a Constraint on Markets,” *Journal of Economic Perspectives*, 2007, 21 (3), 37–58.
- , “Deferred acceptance algorithms: History, theory, practice, and open questions,” *International Journal of Game Theory*, 2008, 36, 537–569.
- Roth, Alvin E.**, “Marketplaces, markets, and market design,” *American Economic Review*, 2018, 108 (7), 1609–58.
- Roth, Alvin E. and Elliott Peranson**, “The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design,” *American Economic Review*, 1999, 89, 748–780.
- , **Tayfun Sönmez**, and **M. Utku Ünver**, “Kidney Exchange,” *Quarterly Journal of Economics*, 2004, 119, 457–488.
- Rothkopf, Michael H.**, “Thirteen reasons why the Vickrey-Clarke-Groves process is not practical,” *Operations Research*, 2007, 55 (2), 191–197.
- Salzman, James, Genevieve Bennett, Nathaniel Carroll, Allie Goldstein, and Michael Jenkins**, “The global status and trends of Payments for Ecosystem Services,” *Nature Sustainability*, 2018, 1 (3), 136.
- Satterthwaite, Mark Allen**, “Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions,” *Journal of Economic Theory*, 1975, 10 (2), 187–217.
- Schnizler, Björn and Dirk Neumann**, “Combinatorial exchanges for coordinating grid services,” *ACM SIGecom Exchanges*, 2007, 7 (1), 65–68.
- Segal, Ilya**, “The communication requirements of social choice rules and supporting budget sets,” *Journal of Economic Theory*, 2007, 136 (1), 341–378.
- Shakespeare, William**, *The Merchant of Venice* 1600.
- Sipser, Michael**, *Introduction to the Theory of Computation*, 3 ed., Cengage Learning, 2012.
- Sönmez, Tayfun and M. Utku Ünver**, “Course bidding at business schools,” *International Economic Review*, 2010, 51 (1), 99–123.
- Sun, Ning and Zaifu Yang**, “Equilibria and indivisibilities: Gross substitutes and complements,” *Econometrica*, 2006, 74, 1385–1402.
- and ———, “A double-track adjustment process for discrete markets with substitutes and complements,” *Econometrica*, 2009, 77, 933–952.
- Teytelboym, Alexander**, “Gross substitutes and complements: a simple generalization,” *Economics Letters*, 2014, 123 (2), 135–138.

- , “Natural Capital Market Design,” *Oxford Review of Economic Policy*, 2019, 35 (1).
- Varian, Hal R.**, “Position auctions,” *International Journal of Industrial Organization*, 2007, 25 (6), 1163–1178.
- , “Online Ad Auctions,” *American Economic Review*, 2009, 99 (2), 430–34.
- **and Christopher Harris**, “The VCG auction in theory and practice,” *American Economic Review*, 2014, 104 (5), 442–445.
- Vickrey, William**, “Counterspeculation, Auctions and Competitive Sealed Tenders,” *Journal of Finance*, 1961, 16 (1), 8–37.
- Weitzman, Martin L.**, “Is the price system or rationing more effective in getting a commodity to those who need it most?,” *Bell Journal of Economics*, 1977, 8 (2), 517–524.
- Will, Pierre-Étienne**, “Creation, Conflict, and Routinization: The Appointment of Officials by Drawing Lots, 1594-1700,” in Paolo Santagelo, ed., *Ming Qing Yanjiu*, Università degli Studi di Napoli, 2002.
- Williams, Steven R.**, “A characterization of efficient, Bayesian incentive compatible mechanisms,” *Economic Theory*, 1999, 14, 155–180.
- Wilson, Robert**, “Incentive efficiency of double auctions,” *Econometrica: Journal of the Econometric Society*, 1985, 53 (5), 1101–1115.
- Wunder, Sven, R. Brouwer, S. Engel, Driss Ezzine de Blas, R. Muradian, Unai Pascual, and R. Pinto**, “From principles to practice in paying for nature’s services,” *Nature Sustainability*, 2018, 1 (3), 145.
- Yokoo, Makoto, Yuko Sakurai, and Shigeo Matsubara**, “The effect of false-name bids in combinatorial auctions: New fraud in Internet auctions,” *Games and Economic Behavior*, 2004, 46 (1), 174–188.