ANNUAL REVIEWS

Annual Review of Economics

Auction Market Design: Recent Innovations

Paul Milgrom

Department of Economics, Stanford University, Stanford, California 94305, USA; email: milgrom@stanford.edu

Annu. Rev. Econ. 2019. 11:383–405

The *Annual Review of Economics* is online at economics.annualreviews.org

https://doi.org/10.1146/annurev-economics-080218-025818

JEL codes: D44, D47

Keywords

financial markets, Internet advertising, electricity auctions, spectrum auctions, cryptocurrencies, combinatorial procurement auctions

Abstract

Market design applies economic principles to the often messy problems of real-world exchange in which goods may not be homogeneous, the identities of trading partners may matter, contracts may not be executed, and even the formulation of trade as balancing supply and demand may be unhelpful. This article recounts the mostly academic research advancing the analysis and design of such markets. Among the highlighted applications are ones involving financial markets, Internet advertising, electricity auctions, spectrum auctions, cryptocurrencies, and combinatorial procurements.

INTRODUCTION

Introductory economics courses begin to teach about markets by focusing on trades of a single homogeneous good with the following three characteristics: The identities of the buyer and seller do not matter, prices equilibrate supply and demand, and a law of one price applies. Real markets are rarely so simple and may be designed to deal with the ways that these characteristics may fail: The identities of trading partners matter, prices play little or no role in market clearing, or prices vary greatly among similar transactions.

In the marriage market, for example, men and women care deeply about the identities of their partners, even when there are no dowries or bride-prices involved (Roth 2016). Prices play little role in limiting the demand for colleges or in guiding the matching of children to schools or kidney donors to kidney patients. There is no single price for Internet advertising. In markets for advertising on Internet search pages, auctions determine a separate price for each ad impression.

The law of one price is tricky to interpret in markets that feature fine distinctions among similar-seeming goods. In electrical power markets, energy is distinguished by the time and location at which it is to be made available: Power in Boston cannot serve customers in Los Angeles, and power at midnight cannot run air conditioning units on a hot midafternoon. In these markets, shorter time intervals and smaller geographic areas have become increasingly common. On the Internet, the most relevant ad to show to a user varies not only across users but also over time for a single user. It can vary depending, for example, on whether the user is shopping to replace a broken kitchen appliance, refinance a mortgage, or plan a vacation.

Another way that the law of one price can fail is that the prices for individual goods may depend on what else the buyer purchases. A consumer who buys a new mobile phone may be offered a discounted price for accessories such as headphones or wireless chargers, so that the discounted price is really for the package of goods. In electricity auctions, a bidder that has fired up its plant to produce some amount of energy may offer additional units of energy at a lower price.

In textbook treatments of market clearing, products are homogeneous, and the only important constraint to be satisfied is to set the quantity supplied equal to the quantity demanded, but those treatments ignore the challenges resulting from heterogeneity among products. Some challenges arise even for a seemingly homogeneous commodity like coffee beans, because grading standards must be set to define product categories. Even within a single grade, not everyone agrees that Kenyan and Rwandan coffee beans are exactly the same. Sometimes, heterogeneity is so important that the standard formulation of the problem in terms of resource constraints and prices obfuscates the deeper economic issues. For example, consider the challenge of assigning television stations to broadcast channels to be used for over-the-air broadcasting. A common way to represent the constraints on this assignment is not as limits on a set of resources to which prices can be assigned but as a graph in which each TV station is a node and two stations are connected by an arc if they cannot be assigned to the same channel without creating interference (Milgrom 2017). A feasible allocation assigns channels to stations so that no two stations that are connected in the graph get the same channel. If we substitute the word color for channel, this is a classic graph-coloring problem: Find a way to color the nodes of a graph so that no two connected nodes are the same color. Graph coloring problems are computationally challenging and cannot be efficiently solved by finding prices to assign to each of the many constraints. Organizing a market to rely on prices to guide these transactions is unlikely to generate desirable outcomes.

After treating markets for a single homogeneous good, textbooks often turn to study substitutes and complements. When goods are substitutes, it has long been known (Arrow et al. 1959) that prices that are adjusted separately in each market can promote simultaneous clearing in all markets, but complementary goods lead to much more complex challenges, particularly if the goods are tightly connected. For example, a property developer who buys land from small landholders for a new shopping center may find that each plot is much more valuable if the adjacent plots can also be acquired, so it may not make sense to approach each transaction separately. McAdams (2018) offers the interesting example in the case of water markets. A certain farm that is using water to irrigate its fields is situated upstream of a nesting habitat for an endangered bird species, which is valued by an environmental group. A downstream farm probably has lower value for water rights, but its use of water preserves the nesting habitat. In this example, if water use and water flow are both marketed products, then the upstream farmer should sell their rights only if the sum of the water-use value to the downstream farmer plus the water-flow value to the environmental group is high.

Most textbook studies of markets focus on resource allocation, with prices playing a supporting role. In financial markets, however, the emphasis shifts from resource allocation to the prices themselves. Traders in these markets seek to anticipate and take advantage of others' trading activity or to defend themselves against such behavior from other traders. Successful trading platforms need to be designed to meet the needs of market participants.

As these examples suggest, the field of market design poses a rich set of problems. Market operators need to structure their trading platforms so that participants can be matched with the right partners, trust and quality can be assured, payments can be facilitated, and prices can be protected against manipulators, and they need to find ways to charge for their services that the matching partners cannot avoid once introductions have been made. As more transactions become automated and move to the Internet, the rules governing transactions have become more explicit, making them easier for academics to study and for bad actors to exploit. Given its practical importance, market design has attracted increasing attention not just in academic journals but also in industry journals, and not just from economists but also from operations researchers, industry professionals, and regulators. Given the recency of the developments, understanding is evolving rapidly, and many research agendas are still underway.

This review seeks to highlight some of the questions and challenges found in the literature about modern marketplaces and is organized around some of the important applications of market design. The following sections study Internet advertising, radio spectrum auctions, electricity auctions, financial markets and electronic trading, cryptocurrencies, procurement auctions, and combinatorial auctions.

INTERNET ADVERTISING

Before 2000, the Internet was mostly regarded as just another advertising medium, similar to print, radio, or television: Internet publishers sold to advertisers under contracts, with prices largely dependent on the number of impressions the site could deliver to a particular category of users. By 2017, the volume of advertising on the Internet had grown to about $88 billion, with sponsored-search and display ads each accounting for about 45% (PricewaterhouseCoopers 2018).

The first big changes to distinguish advertising on the Internet came with sponsored search, according to which an advertiser paid to have its ads appear alongside organic search results when a user searched for particular keywords. For example, an auto insurance company might select "insurance" as one of its keywords, triggering its ad whenever a user searched using that term. Unlike traditional print and broadcast advertising, for which ads might be shown to mostly uninterested users, search ads would more often be shown to users who were interested in the advertiser's product or service. Google and other search engines soon introduced auction systems that, each time a user searched, would determine which ad to show and what price the advertiser would pay. With billions of searches being run every day, each taking just a few milliseconds to complete, Google was soon running far more auctions than any other company (Evans 2009).

Sponsored searches quickly began to generate huge revenues for Google, and that inspired other Internet publishers to investigate ways to target users on their websites as well. With no search terms to rely upon, they began to gather as much other information about each user as they could and used it to target advertisements. A user who had bought a flight from Chicago to New York City or read articles about new plays on Broadway might soon after be shown ads for Manhattan hotels, restaurants, and attractions. Publishers and others developed systematic ways to report and share information so a user who bought a ticket on an airline's website could see a travel-related ad on a completely unrelated website. Initially, advertisers and publishers depended on traditional advertising contracts to govern their relationships, but as economic exchange technologies improved, ad exchanges began to run auctions both to select which ad to post for each opportunity and to set a separate price for each.

In order to expand their businesses, there were two large and novel trust challenges that the new Internet-based ad markets needed to address. In the past, ads on a television show or in a newspaper were priced in proportion to the number of ad impressions, called per-impression prices, so an advertiser might pay $15 per thousand impressions in certain media. To check that the ads were actually shown, it was easy to turn on the TV or buy a newspaper, and independent sources could verify claims about circulation or viewership numbers. On the Internet, however, with individually targeted ads, how could an advertiser verify, say, Yahoo's claim that it had shown 1.3 million ads to various individuals on a diverse set of its web pages? Furthermore, when an ad for soccer balls appears in *Sports Illustrated* magazine, the advertiser can be pretty confident about the demographics of its audience and can estimate the ad value from that. But, how can the advertiser trust that its Internet ads were shown to users who might be interested in its product or service?

Search engines tackled these novel challenges by charging advertisers not for the impressions they bought but for the clicks on their ads (Lahaie et al. 2007). An advertiser still cannot monitor where or how often its ads are shown, but it can count the visits to its website originating from any source and can avoid paying to show ads to users who are too disinterested to click on the ad. This pricing solution was a practical economic innovation to mitigate the trust problem. The solution is not a perfect one: Some clicks may be fraudulent ones by dishonest publishers or the bots of a firm that wishes to increase a competitor's costs. Nevertheless, click-based pricing was widely adopted as an improvement, and it inspired other changes in the market as well. Search engines like Google, which were paid for clicks on the ads they showed, invested in research predicting the click rates of different ads.

Another innovation that search engines developed was broad match, which improves matching between advertisers and impressions by making it easier and less costly to participate in the relevant auctions (Amaldoss et al. 2016; see also Dhangwatnotai 2011). To illustrate broad match, suppose that a company specifies property insurance as a keyword and that the user searches for something else, such as fire insurance or insurance protection against property damage. Although neither of these searches includes the exact phrase property insurance, the insurer might still want to show its ad to these users. Broad match allows the search company's algorithms to judge which of these phrases should trigger the ad, potentially allowing more effective ad campaigns. Moreover, the search engine has an incentive to predict the click rate of these broad matches accurately because improving the match between ads and opportunities leads to higher ad revenue.

Eliaz & Spiegler (2016) further develop these ideas, showing that a properly designed broad match algorithm improves consumer search and increases competition among advertisers and can improve competition among publishers as well. Broad match was the first of several technologies that simplify bidding, allowing advertisers to express values simply for one narrow set of outcomes, while relying on search providers' algorithms to infer values for other outcomes and promote those. The market design literature has paid only limited attention to simplification, although Levin & Milgrom (2010) provide an exception, studying how conflation—the suppressing of distinctions—can sometimes simplify and lubricate the operation of markets.

When several advertisers compete to show their sponsored ads alongside the natural search results, the search provider conducts an auction to decide which ad to show. In the early history of these auctions, winners paid the prices that they had bid. Two factors combined to make that system problematic. First, there were multiple ad spots on each search page. Second, bidding was often controlled by automated bots, which could observe the bids from previous auctions for the same keyword. The bots were often programmed to predict that the same bids would be repeated and would optimize their own bids accordingly. The economic effect of the bots operating in such an environment generates a characteristic dynamic pattern of bids and prices. To illustrate this pattern most simply, let us suppose that the number of bidders is two, the number of slots is two, and there is a low reserve price for each slot. In a sequence of auctions, the losing bidder is inclined to increase its bid by just enough to win the first position. Prices climb, one penny at a time, until one bidder concludes it would do better to reduce its bid to win the second position at a low price, rather than to increase its bid to win the first position at a high price. Once the low bidder does that, however, the high bidder finds that it can still win first position with a much lower bid. It reduces its bid to be just enough to win and the whole cycle starts again. These cycles, known to economists as Edgeworth cycles, are described and graphed by Edelman & Ostrovsky (2007).

This instability of prices and winners can be a problem for many reasons, not least being that the resulting allocations are inefficient because the highest-value bidder wins only about half the time. Google responded by changing its auction rules. Instead of paying the amount of its winning bid, a winner would pay only the minimum price it would have needed to bid to win its position. So, if one bidder bids \$2 and another bids \$1, the \$2 bidder would still win the top position but pay a price of just \$1.01 because that is the lowest price that would win the first position. This auction format is now widely known as the generalized second-price auction. Adopting this format eliminated bidders' incentives to make frequent bid adjustments and stabilized both the price and the winner in these auctions.

The theory of the generalized second-price auction is studied by Edelman et al. (2007) and Varian (2007) using noncooperative game theoretic models. Among the full-information Nash equilibria of this game is one in undominated strategies at which the prices are competitive equilibrium prices. This was called the locally envy-free equilibrium and the symmetric equilibrium in the two papers, respectively. In this Nash equilibrium, the auction outcome is efficient, and the prices coincide with those of the dominant strategy solution of a Vickrey auction. However, these analyses fail to apply general game theoretic reasoning to select an equilibrium. These same conclusions about prices and efficiency of the pure-strategy equilibrium were later confirmed to hold for two general game theoretic refinements: test-set equilibrium (Milgrom & Mollner 2018a) or extended proper equilibrium (Milgrom & Mollner 2018b).

Athey & Ellison (2011) emphasize that, besides the advertiser and the publisher, there is another important party in an ad auction: the consumer. Relevant ads improve consumer experience, helping them to find the sellers from whom they prefer to buy. Improving consumer experience not only benefits consumers directly but also increases their reliance on search ads and with it the sales and revenues for advertisers and publishers, too.

Sponsored-search advertising, which now accounts for about 45% of all Internet advertising revenue, achieved its remarkable success by targeting ads to the specific and current desires of consumers, as indicated by the search terms they typed. This inspired imitators hoping to target ads on other websites as well, using whatever information they could about a consumer and their current interests. Initially, this so-called display advertising was handicapped as many advertisers, worried about the lack of good information to guide their ad placements, participated in auctions only for sponsored-search advertising or for search plus display ads on a few carefully selected websites. With so many opportunities to advertise on the web and such limited participation by advertisers, the early display advertising markets were thin and included many low-value matches. Limited information was harming this market.

There have been several attempts to solve this information problem. Hu et al. (2015) describe the use of pay-for-performance advertising in which advertisers pay not for clicks but for some measure of user behavior after the user has clicked on an ad. An early version of this idea introduced by Google in 2004 was its Smart Pricing program, which was designed to allow Google to use bids for sponsored-search advertising for ads shown on other websites. A participating advertiser in this program specifies what it means by the performance of a user's click. For one advertiser, it might mean that the click resulted in a sale; for another, it might mean the user completed a form or visited a second page on the advertiser's website. Under the Smart Pricing contract, the advertiser allows Google to monitor the performance of clicks on its ads on the advertiser's website. Smart Pricing then adjusts an advertiser's sponsored-search bid per click on each different website based on the predicted relative performance of clicks from that site. This makes it cheaper and easier for a sponsored-search advertiser to expand its advertising program onto the wider Internet.

Smart Pricing is a response to a traditional challenge for markets: the problem of adverse selection. Ad opportunities are very heterogeneous, with some impressions presenting a much greater likelihood of a large and profitable sale. For example, a click on an ad for mortgage financing is much more likely to be shown to a serious potential borrower if the ad is on a financial website compared to a sporting goods website. Advertisers who are less sophisticated than others in picking among impressions are likely to get an unfavorable or adverse selection of impressions and suffer a lower yield on their advertising budgets. Smart Pricing is one way to mitigate that by allowing the search company to adjust the bid for each ad impression based on a statistical assessment of the likelihood that an ad will result in a good outcome.

Another attempt to deal with a part of the adverse selection problem adjusts the auction rule that might be used by a large publisher that shows ads under contract with its regular contract advertisers as well as to others. An example of a contract advertiser might be a shopping mall that is advertising its new weekend hours both online and in a newspaper. Its ads are shown to large numbers of people in its geographic area, intending to create awareness among them without inspiring any user to click on an ad. Even if the mall's promotions are successful in attracting customers, it will find it difficult to distinguish which ads or promotion efforts are most responsible for its success. The publisher's other advertisers might be sellers more narrowly targeting customers who are thought to be particularly good sales prospects, for example, because they shop frequently in the online stores, have shown interest online in their products, have high incomes, or have responded well to past online advertising. Advertisers in this second category are sometimes called performance advertisers because they typically seek some specific online performance from the consumer such as buying a product, viewing a video, or visiting an online store. These performance advertisers seek out better information about the individual online consumers, which creates a potential for adverse selection. To the extent that the targeted consumers are people with higher incomes or who are more likely to respond to online advertising, their selection by performance advertisers detracts from the audience left for the contract advertisers.

Arnosti et al. (2016) investigate an auction design that aims to mitigate the adverse selection against contract advertisers suffering an information disadvantage compared to performance advertisers. In their theoretical model, a publisher has committed to sell some fixed number (or fraction) of its relevant impressions to some contract advertiser that can anticipate the average value of a random impression in that set but cannot observe its profits from individual impressions. The model also has two or more performance advertisers, each of which knows its own value for any individual impression. In the second-price auction traditionally used for allocating Internet display ad impressions, there is good matching for performance advertisers but adverse selection against the contract advertiser. The paper studies whether there is an alternative auction design that might have better properties. It formulates the problem by assuming that each individual advertiser's value is the product of an idiosyncratic match value and a common value that depends on the user.

The paper deviates from the usual mechanism design approach that is more common in the auctions literature. Instead of solving a maximization, it approaches the problem axiomatically, seeking to devise an auction design with five good properties: The auction should be (*a*) strategy-proof (so bidding truthfully is a dominant strategy), (*b*) efficient in its assignments among performance advertisers, (*c*) anonymous (symmetric among performance advertisers), (*d*) false-name proof (no bidder can reduce its price by submitting an extra bid, and the seller cannot raise its price by submitting a high shill bid), and (*e*) adverse-selection free (the contract advertiser's distribution of match values must match those of the total population). Setting aside the second property, none of the other properties are implied by efficiency, and none of the five are implied by revenue maximization, so there is no guarantee that such an auction, even if it exists, would always perform well. The paper shows that there is a unique family of modified second-bid auctions that has the five properties. The family is parameterized by a single parameter $\alpha \geq 1$. In these auctions, the highest bidder wins if the ratio of its bid to the second highest bid strictly exceeds α, and in that case the impression is awarded to the highest bidder at a price equal to α times the second highest bid. Otherwise, the impression is awarded to the contract bidder.

To assess the performance of this novel design, the paper includes a numerical analysis of its performance assuming that the match values are drawn independently from a power law distribution. The power law describes a special environment in which most of the expected value of good matching comes from very good matches. Assuming that there is a contract that specifies the fraction of impressions that must be assigned to the contract advertiser, they find that the worst-case ratio of the value achieved by the mechanism to the value of the full-information optimal matching is nearly 95%.

The previous analysis takes the set of contract ads as fixed. How should a publisher decide the allocation between contract ads and performance ads? Sayedi (2018) studies that question, taking the perspective that contract ads can set a reserve price in an ad auction. According to classical auction theory, reserves are most valuable when the number of likely bidders is small. Setting the right fraction of contract ads allows the seller to ensure that carefully targeted advertising, which delivers high value to performance ad buyers, can also maximize revenues for the seller.

The Internet advertising market is huge and quite dynamic, as new targeting methods, new ways to include online advertising in larger ad campaigns, and new options for how even simple search ads are shown continue to be developed.

One recent development in sponsored-search advertising is variable-size ads. In the first decade of sponsored-search advertising, all ads were restricted to be the same size, consisting of a few lines of carefully chosen prose and a link to some website. Today, sponsored-search ads can vary in size, from perhaps 3 to 18 lines. When all ads were the same size, one could conceptualize the problem as which ads to show from the limited number that fit on a page and which ad to show in first position, second position, and so on. With variable-size ads, the number that can fit on a page is not fixed, the position number does not fully determine the location of the ad on the search page, and the size of an advertiser's ad may depend on the other winning bids. This created new challenges in determining which ads to show and how much to charge and in deciding how bids should be structured.

Some initial attempts to deal with this had bidders offer different prices for ads of different sizes. Even so, fitting ads onto a page to maximize total value is what operations researchers call the knapsack problem, which is a computationally difficult problem (i.e., it is NP-hard). Pricing for such problems is correspondingly hard. If the number of available lines is small, this might conceivably be solvable by using a Vickrey auction, but such auctions can sometimes lead to very low prices, even in highly competitive situations (Ausubel & Milgrom 2006). No consensus has yet emerged about the best way to solve this problem.

The market for Internet display advertising received another shock in 2016 with the advent of header bidding (Wang 2018). Prior to that time, a publisher who had space to sell on its page would send its opportunity to an ad exchange, where qualified advertisers would bid on impressions. If the publisher got back an acceptable price, it would show the impression; otherwise, it might pass the impression along to another exchange or show its own house ad. Now, however, many publishers are soliciting and comparing bids from various ad exchanges. This is a deeply problematic market organization. For example, suppose that the bids in the first exchange are $5 and $3, while those in the second exchange are $10 and $2. Using a second-price auction, the price in the first exchange is $3 while that in the second exchange is $2, so the bidder who bid $5 wins over the bidder who bid $10. This organization creates a huge inefficiency, and there are other problems, too. The bidders who bid $5 in one exchange and $10 in the other could be the same bidder, running two different kinds of ad campaigns and unaware that it is bidding against itself. This pattern is most unlikely to persist. It could lead advertisers to prefer exchanges that use first-price auctions and, depending on what other services the exchanges offer, it could then lead to further consolidation in the industry as each bidder seeks to use just one exchange for all its ad campaigns.

RADIO SPECTRUM AUCTIONS

Until the 1990s, radio spectrum licenses were nearly always allocated based on administrative procedures to determine the public interest, colloquially known as beauty contests. Building on work by law student Leo Herzel (1951), Ronald Coase (1959) became the most famous early advocate of using auctions instead, but his recommendations long fell on deaf ears. Critics mocked the idea, saying that the chances of using auctions to allocate radio spectrum licenses in the United States were about the same as those of the Easter bunny winning the Preakness. But in 1994, Coase's fantasy became a reality, and one that was soon widely copied around the world.

Why did it take so long? Coase's early analysis was rooted in the traditional textbook theory of markets, which had been applied successfully for Treasury securities but not much else. Because Coase's analysis gave no consideration to the many complexities cited in the introduction to this paper, it was ill-suited to guide the creation of the auctions for multiple, heterogeneous licenses that were needed in the United States. Before 1994, failure to account for those complexities had resulted in a series of failed auctions in other countries, as described by McMillan (1994).

In 1993, when Congress first authorized the use of auctions for radio spectrum in the United States, the Federal Communications Commission (FCC), having no previous experience with auctions, put its economists in charge. The first big auction would entail selling more than a thousand different radio spectrum licenses, all for use in mobile phones, but distinguished by the geographic areas they covered and the frequencies to be used. FCC economist Evan Kwerel led this celebrated effort and issued a proposal that extensively referenced academic work on auction theory. After a months-long process of hearings and debate, the FCC adopted what was called the simultaneous multiple round auction (SMRA). Eventually, that new auction design was used for more than $100 billion of radio spectrum sales around the world (Milgrom 2004).

The SMRA proceeds in a series of rounds, with bidders free to place bids on many licenses, provided that they exceed the previous highest bid by some minimum amount. The design also included the Milgrom-Wilson activity rule, which specified roughly that no bidder could increase its activity from round to round, which meant that it could not bid on a larger volume of licenses than in the previous round. Adapting the logic of Kelso & Crawford (1982) to this application, Milgrom (2000) shows that if licenses are substitutes for each bidder and bidders bid for the most profitable package at the lowest possible price in each round, then the eventual allocation would be nearly efficient (to within a price increment), and the final prices would be approximately competitive, market-clearing prices. The activity rule, which ensures that the auction proceeds to its logical conclusion at a reasonable pace, does not alter that theoretical conclusion.

Much was learned from the early uses of the SMRA design. One problem of the design is that it is slow and can take many rounds. When the auction involves the sale of multiple units of several kinds of products, essentially the same process can be accelerated by having the auctioneer, rather than the bidders, name prices and set a single price for each type of good, with the price being increased in each round for any category in which demand exceeds supply. This can substantially speed up the auction process (Milgrom et al. 2012).

Another problem is that bidders may try to collude or divide markets. For example, in the auction for third-generation mobile licenses in Germany in 1999, Mannesmann and T-Mobile each managed to win an equal number of licenses without competing against one another. Citing that case and others, Klemperer (2002, p. 170) argues that "what really matters in auction design" is traditional industrial policy to prevent "collusive, entry-deterring and predatory behavior." The importance of these elements is also endorsed by other analysts. Cramton & Schwartz (2002) find that bidders sometimes collude by using bids to signal, that is, to make threats and promises and to offer deals. Ausubel et al. (2014) find that bidders in an SMRA will often find it profitable to exercise market power by reducing demand to reduce the prices that they pay.

These problems are similar to ones found in other concentrated markets, and the impact on prices and efficiency can be very large. For example, suppose there are two items for sale and two bidders who have values of \$5 and \$3 per item. The efficient outcome would assign both items to the first bidder for a total value of \$10. Suppose both bidders submit bids in the form of demand functions, that is, maximum prices for one item or two and that the price is set as the lowest market-clearing price. There is a perfect equilibrium of this game in which each bidder bids its full value but for only one item, each wins one item, and the market-clearing price is zero. If the prices are determined dynamically with the bidders taking turns responding to prices announced by the auctioneer, who starts at zero and raises the price by one if there is excess demand, then this outcome is the unique subgame perfect equilibrium. SMRA designs are vulnerable to the usual problems of collusion in uniform price auctions.

Yet another early concern expressed about the SMRA was that it could not deal effectively with the reality that licenses covering different regions in the United States are complements, not substitutes (Charles River Associates & Market Design, Inc. 1997). Mobile service companies in the United States have often competed for customers by claiming to have the best nationwide coverage, perhaps claiming that a customer from Boston who drives west to San Francisco will find that her phone works everywhere along the way. For such a company, acquiring licenses of some kind in Boston is not so valuable unless it also acquires similar licenses in Chicago and San Francisco. When the value of a collection is more than the sum of the individual values, that is a sufficient condition for the licenses to be complements. A bidder in an SMRA is exposed to the risk that it might win some of the licenses it seeks only to find that the prices are too high for the remaining licenses needed to sustain a viable mobile network. That is a bidder's exposure problem, and laboratory experiments (Ledyard et al. 1994) began to suggest that auction designs that avoided that problem might result in higher efficiency outcomes than the SMRA.

Auctions in which bids applied not just to individual items but to complete packages or combinations of items are called package auctions or combinatorial auctions. In the example above, a package auction might allow a bidder to bid for licenses covering certain frequencies nationwide or in all the major cities. The FCC's interest in package auction research led to a book entitled *Combinatorial Auctions*, edited by Cramton et al. (2006). An article in that book by Ausubel et al. (2006) introduced a new auction format, now known as the combinatorial clock auction (CCA), which with modifications has become one of the most common auction formats used for selling spectrum worldwide. It was designed to eliminate both the collusion problem and the exposure problem.

The CCA is a particular package auction that works in two main stages. In its clock stage, it functions much like the clock auction above, in which prices are increased for categories of items for which demand exceeds supply. However, the bids made in this clock stage are package bids; that is, a bidder is offering only to buy the whole package at the specified price. When the clock stage of the CCA is over, there is a supplementary stage in which bidders can make additional package bids. Each bidder's bids in the supplementary stage need to be consistent with the bids made in the clock stage, but the precise consistency requirement has varied from auction to auction. When all the bids are finalized, the winning bid is the set of bids that maximizes the total bid price, subject to the constraint that the item assignment is feasible. In the CCA, the prices that winning bidders pay are not equal to the bids but are instead determined by a core-selecting pricing formula, according to the principle suggested by Day & Raghavan (2007), Day & Milgrom (2008), and Day & Cramton (2012). The resulting prices are typically close or identical to Vickrey prices. In particular, each bidder's bids have little or no impact on the prices it pays for what it wins. To the extent that this impact is approximately zero, there is nearly zero incentive for individual demand reductions to reduce prices, making it much harder for bidders to divide markets in the auction alone. However, the Vickrey auction does open the possibility of joint deviations that are profitable for all participants. These might be supported by other interactions among the bidders. Indeed, to the extent that the auction is part of a larger set of interactions among the bidders in which the bidders can compensate or punish one another for bad behavior in the auction, any analysis that focuses on the auction in isolation might fail to characterize the most likely bidder behavior.

A recent book entitled *Handbook of Spectrum Auction Design* (Bichler & Goeree 2017) reprints many of the articles on this topic, including sections on the SMRA, the CCA, alternative auction designs, and secondary markets.

The Incentive Auction

As wireless broadband services have grown in importance, all the available frequencies in the usual mid- and low-frequency ranges have been assigned for some use. To facilitate the continued growth of wireless broadband, some method was required to reallocate frequencies, and in the United States, that is easier to do when there is a plan to compensate those who will lose access to the spectrum. In 2012, the US Congress changed the law to permit the FCC to conduct a broadcast incentive auction (or just incentive auction) to buy certain TV broadcast rights from station owners, move other broadcasters to new channels while compensating their retuning costs, and sell licenses to use the cleared frequencies for other uses.

What made this transaction especially difficult is that what station owners would sell (TV broadcast licenses) was not in any simple, proportional correspondence with what broadband service companies would buy. Even checking whether there is a feasible way to assign channels to stations without creating interference is mostly representable as a graph-coloring problem, as described in the introduction. NP-complete problems at this scale are large enough that no known algorithm can solve all such problems in reasonable time, despite the enormous advances in modern computers, algorithms, and software. In practice, optimization among such assignments is even harder. This computational challenge implies that Vickrey prices cannot be reliably computed, since the Vickrey price of any TV station is the difference between the optimal value of the assignment with the TV station on or off the air. Even approximate optimizations that achieve 99% of the optimum are not sufficient for practical Vickrey pricing in a setting with more than 2,000 stations. Since 1% of 2,000 is 20, an error of that size in one computation error with perfect computation in the other would lead to an estimated Vickrey price for a station that misses the correct value by an amount equal to the value of 20 average stations. Errors smaller than that could not be guaranteed in these computations, and pricing errors of 20 times the value are as unacceptable to the bidders as they are to the auctioneer.

The Vickrey auction has other serious defects for this application. An important one is that every bidder in a Vickrey auction would be asked to report the value of its station and trust the auctioneer to keep the report secret and compute the winners and prices correctly according to the rules. Even if the computations could be performed perfectly accurately, many bidders in the auction would likely be unconvinced that the government could be trusted to do all that. Distrust could undermine a bidder's incentive to bid truthfully or possibly even to participate. The actual auction design was obviously strategy proof in the sense of Li (2017). According to that concept, even if the bidder does not believe that the computations are done accurately or honestly, it can still do no better than to bid straightforwardly according to its station value.

Another big problem of the Vickrey auction is that it could be very expensive. Simulations showed that a simple Vickrey auction, even if feasible, might have paid prices to low-powered stations serving few customers that were as high as those paid to high-powered stations in nearby geographic areas. It is as if the money were a bribe paid to a station to prevent interference, rather than compensation for its loss in going off the air.

An additional problem with the Vickrey auction is that it does not respect the budget constraint that the total price paid to the sellers cannot exceed the revenue from selling the resulting licenses. Yet another is that it does not preserve winner privacy, meaning that others would be able to learn something about what the winners have earned. Finally, the Vickrey auction is not group strategy–proof, meaning that there can be ways for bidders to deviate jointly to increase the profits of all bidders in the set.

None of these defects are shared by the actual auction design adopted by the FCC (Milgrom & Segal 2018), which is a new kind of descending clock auction that offers each station a different price depending on its own characteristics. Still, the Vickrey auction is the unique strategy-proof auction that guarantees efficient outcomes, so a fair comparison requires estimating the loss of efficiency from the descending clock design. Using estimated station values and simulating the auction results using stations in the region around New York City, Leyton-Brown et al. (2017) compare the results of the FCC's descending clock auction to those of the Vickrey auction. In those simulations, the descending clock auction, on average, adds about 5% to the minimum value of stations removed from the air (a direct efficiency loss); it also reduces the cost of the procurement by about 24%. Thus, giving up exact optimization led to considerable cost savings.

ELECTRICITY AUCTIONS

Electricity markets have posed a series of specialized and unique challenges for regulators over the years. Some of these challenges are related to variations in supply and the characteristic supply functions for plants, which include substantial and widely varying fixed and marginal costs for different power generation technologies. As a matter of efficiency, the plants with the highest marginal cost are effectively standby capacity; that is, they are only used in periods of unusually high demand. The interconnectedness of the electricity system, however, implies that a failure due to inadequate capacity is suffered by all generators and customers. System reliability is therefore a public goods problem, making it unclear whether anyone would voluntarily hold standby capacity and how such capacity should be compensated. Another more commonplace issue in these markets is market concentration, which is particularly salient in periods of peak demand. Capacity is a special problem for electrical power markets because, historically, demand has been very inelastic, so even tiny shortages can lead to huge price fluctuations to bring supply and demand into balance. Technical issues, such as the surprising implications of Kirchoff's laws governing electricity flows, loom large in some analyses, particularly ones that call for changes in or upgrades to power transmission networks.

Even setting aside all the characteristics that make electricity markets special, which can and do fill their own journals (a leading one being *The Electricity Journal*), these markets also emphasize some of the specific themes highlighted in the introduction. On a hot summer afternoon when air conditioners are running, home demand for electricity may be many times higher than on a cool night when consumers are sleeping, and the marginal cost of electricity can vary dramatically with usage. Although consumers have traditionally measured electricity use as an aggregate quantity of kilowatt hours over a long period such as a month, economists cannot treat power as a single commodity with a single price. One of the main challenges for power markets is to balance supply and demand in the system hour by hour and even minute by minute over the course of each day. Joskow & Wolfram (2012) emphasize that the lack of real-time metering in the past has meant that consumer prices cannot reflect the huge variations in marginal costs during the day, but they also highlight how this is on the verge of changing. Advances in real-time metering can enable dynamic pricing and lead to important improvements in system operations, if only consumers can be given appliances that allow them to shift demand to low-price times of day. A second challenge for these systems is to provide adequate long-run supply, known as capacity, to these systems. Power storage may be another part of the answer. Attempts to produce efficient batteries to smooth the variation in supply from natural sources like sun or wind may make it possible to use stored energy from low-cost sources during periods of high demand.

Many papers on electricity auctions treat the problem of designing an auction market to promote efficient supply and competitive pricing as theoretically similar to designing an auction for Treasury bills or other securities. Focusing on short time periods, Fabra et al. (2006) study the performance of different auction rules for energy markets in the presence of market power. The paper assumes that short-run demand is unresponsive to prices, so the quantity of energy to be purchased is exogenous. It contrasts uniform price auctions, in which all suppliers receive the same price for energy, with discriminatory auctions, in which the winner's energy price is determined by its bid. Studying an undominated Nash equilibrium in a full information duopoly model, the paper finds that the discriminatory auction format leads to lower costs but that the efficiency comparison is ambiguous.

However, energy auctions are trickier than financial auctions because nonconvexities in power generation are important. There can be large fixed costs in turning on and ramping up a generator to supply power in the mid-afternoon and ramping it down when service is no longer needed (Wilson 2002). Standard price theory analyses do not dig deeply into these sorts of details, and auctions that are designed to operate independently in setting separate prices for different products may fail to coordinate producer activities well.

Cramton (2013) dives deeply into the critical and difficult relation between spot energy markets and available capacity. Until there is widespread real-time metering with supporting home appliances to allow demand shifting over time, demand will remain inelastic, so efficient energy use throughout the day cannot be guided by a system of energy prices alone. On a hot day, energy overuse can lead to brownouts and blackouts, which can cause huge losses of human welfare in terms of health, comfort, and production. With such unresponsive demand, an efficient supply system requires tremendous standby generation capacity in which some generators are used during just the highest 1% of demand hours and are inactive the rest of the time. Such generators have sometimes been compensated for by having enormous spikes in energy rates during the highest demand hours, but such a system has multiple disadvantages. One is that it encourages market manipulations by suppliers who might forecast critical times to schedule plant maintenance, creating shortages to drive up prices. Another is that critical periods must actually occur periodically at the right rate to finance capacity additions. Since it is difficult to forecast rare events based on history, it is unlikely that anyone could regulate a system successfully in that way. Cramton suggests that one solution may be to devise a system that pays separately for capacity. Another is for the regulator to buy options contracts to purchase energy from suppliers at a given strike price, using the market's forecast to replace the regulator's forecast. This proposal would reduce the revenue and cost fluctuations for both producers and any public energy purchaser and limit suppliers' abilities to benefit from critical energy shortages.

Maurer & Barroso (2011) are notable for providing a comprehensive list of case studies of capacity auctions and energy auctions. These highlight the variety of solutions that regulators have tried to address the capacity problem.

Cramton (2017) highlights the structural changes in modern power markets. Regulators have been trying to encourage the use of renewable energy, better battery storage technologies, and smart home technologies, including both time-of-day metering and devices that can respond to and schedule activities according to those prices. Renewable energy is often favored for its environmental benefits quite separately from its impact on the energy markets, but by adding capacity to markets with fixed demand, new renewable sources can significantly depress prices paid to suppliers. Both battery storage technologies that allow power produced at one time to be supplied at another and smart home devices that allow demand to respond to prices can enhance efficiency and improve system reliability. As for almost every application of market design principles, designers of electricity markets need detailed knowledge of the setting and the relevant technologies if they are to design markets that stimulate reliable electricity supply, adequate investment in capacity, reduced emissions, and low costs.

FINANCIAL MARKETS AND ELECTRONIC TRADING

The emergence of automated, electronic trading has led to important changes in the operation of financial markets. Bids and offers are now collected in an electronic order book—an organization known as a continuous double auction, which is similar to mechanisms that traditionally led to efficient outcomes in many experimental economics laboratories (beginning with Smith 1962, 1965). This market organization, with its very low costs and automated trading, has enabled a new class of high-frequency trading strategies, which aims to take advantage of small, short-lived arbitrage opportunities. These strategies entered the public imagination when they were unmasked in the popular book *Flash Boys: A Wall Street Revolt* (Lewis 2014).

Budish et al. (2015) kicked off the recent research into high-frequency trading. They document a fact that, in retrospect, is entirely unsurprising: It takes time for trading on one exchange to affect related prices on another exchange, and that offers a brief window of opportunity for the arbitrageur who moves fastest to earn a profit. Although this model incorporates no actual asymmetric information, the arbitrage imposes losses on those who post prices just as if the market maker suffered from adverse selection, leading to a spread between the bid and ask price, as observed by Glosten & Milgrom (1985), upon whose model their own is built. The biggest welfare losses in this model emerge as professional traders race to be a few microseconds faster than their competitors to be the first in line when an arbitrage opportunity emerges. The large, rent-seeking investments they make in high-speed communications and processing technologies add no value to the economy.

As a solution, the paper suggests replacing the continuous time order book with what it calls periodic batch double auctions. In this system, the market operator would collect orders as they arrive and periodically (for example, every second) construct demand and supply curves from the orders to determine the clearing price and allocation. When speed races are not going on, this system would in practice nearly replicate the current one. However, it would operate differently in situations where milliseconds matter. Orders occurring within milliseconds of each other would nearly always be cleared in the same auction, eliminating most of the advantage of acting very quickly on public information. In that way, the proposed reorganization would reduce incentives for wasteful investments in speed and make participation safer for the slower traders.

Budish et al. (2018) study the incentive of exchanges to develop rules and the technology to run such batch auctions. In an analysis that is reminiscent of arguments for government-sponsored research, they find that any innovators who incur the costs of developing such a system and overcoming its regulatory hurdles would not be likely to capture much benefit, because the system could not be patented and would be too easy to mimic. Consequently, they argue, the regulators themselves need to take the initiative to make it easy to implement such a system.

In a related development, Biais et al. (2015) study the equilibrium level of investment in fast trading technologies in a rational expectations model and determine that social welfare could be improved by restricting this investment. They argue in favor of the use of Pigovian taxes on investment in the fast trading technology.

Another set of papers studies how to match trades among buyers and sellers over time. Brusco & Jackson (1999) study a setting in which there is a fixed cost incurred to trade across periods. The optimal market design reduces those costs by identifying a single party to serve as market maker, acting as a buyer or seller of last resort to smooth trading over time. Vayanos (1999) studies a multiperiod double auction model in which public information arrives over time and investors have exogenous reasons to adjust their holdings. In this model, it is risky for investors to hold too much or too little of a single security. When a large investor needs to adjust its holdings, it can reduce the adverse price impact by spreading out its trading over time, and at least some such spreading is always optimal. In equilibrium, however, this behavior is self-defeating. By spreading out their trades, investors reduce the short-run liquidity of the market, causing even the small trades of other investors to have price impacts. The opportunity to trade frequently and spread out large trades reduces investor welfare. In the limit, allowing trades to be spread over many trading periods results in a welfare loss of order $1/N$, where N is the number of traders.

Rostek & Weretka (2015) refine the analysis of Vayanos (1999), in which information arrives in the form of publicly observed dividend payments but assets can be traded in multiple periods between those payments, decoupling the arrival of news from the frequency of trade. The paper finds that this decoupling reverses the previous result: Increasing the frequency of trade while holding the frequency of news arrival fixed increases investor welfare. Intuitively, the key difference in such models is that the cost of gradual portfolio adjustment is diminished because the risk of new information moving prices in the interim is diminished. In equilibrium, increasing trade frequency is welfare improving.

Du & Zhu (2017) study a model very similar to that of Rostek & Weretka (2015), which assumes that the frequency of trade does not affect the exogenous public information about the risky asset being traded. In contrast to the two previous models, however, traders' values are assumed to be interdependent, so trading exposes investors to adverse selection. In this context, frequent trading involves a trade-off. High-frequency trading allows investors to react quickly to new information and realize the gains from trade associated with the reallocation of assets. However, increasing opportunities for trade also encourage traders to avoid the cost of their price impact by submitting less aggressive demand schedules, which tend to reduce the beneficial reallocation of assets—an effect that is exacerbated by adverse selection.

In these three papers, large traders break up their orders to reduce their adverse price impact. This tendency would not be changed by the use of batch auctions as proposed by Budish et al. (2015), so there would still be significant scope for sniping in which arbitrage-seeking traders anticipate an investor's future trades and try to trade ahead of it in anticipation that market prices will move. Kyle & Lee (2017) argue that a true continuous time exchange would avoid forcing large traders to break up their orders, investing time and resources executing the associated optimal trading strategy. Instead, that function should be automatic: Trades should be spread continuously at a rate specified by the trader, and trading should be continuous in price, quantity, and time. Kyle & Lee call such a mechanism continuous scaled limit orders. In such a mechanism, messages sent by a trader to the exchange would specify a quantity of the asset to be purchased or sold at prices in the interval $[p_{\mathrm{L}}, p_{\mathrm{H}}]$ at a maximum rate of $U_{\max}$. At each instant in time, the exchange computes the flow demand and supply and determines the market-clearing price. Orders can be removed from the exchange at any time, leaving the trader with the stock position that has accumulated up to that point in time. Under such a mechanism, faster traders do have an advantage in the sense that they can react faster to changes in fundamentals and respond by updating their orders more quickly. However, this advantage is limited by the ability of a slower trader to protect itself by limiting its flow rate of trade, which limits the damage it incurs before it can update its order. This system obviates the need for a trader to expend time and energy determining how to optimally break up its trades (this is done for traders by the exchange itself).

An important advantage of this continuous scaled limit orders design is that it easily accommodates the possibility that a single asset may be traded on multiple exchanges simultaneously and that multiple assets might be traded simultaneously, which in other systems could raise serious issues regarding how market clearing should be synchronized across assets and exchanges. However, this market design solution differs radically from the market-clearing mechanisms used in real-world financial marketplaces, which would make it challenging to implement in practice.

Harris (2013) suggests an alternative market design solution to address problems related to high-frequency trading. Rather than running frequent batch auctions, this paper proposes adding a random delay of 0–10 ms to orders submitted to an exchange. Clearly, such a design will partially mitigate the speed advantage of high-frequency traders, but it might fail if high-frequency traders could simply send multiple messages to the exchange.

Melton (2017) describes the ideal latency floor mechanism, introduced by the major electronic trading venue Thomson Reuters Matching, as a means of addressing concerns regarding high-frequency trading. That mechanism intercepts messages to the limit order book and randomly buffers them, so that when the messages are released from the buffer they generally reach the limit order book in a different temporal order. While this is conceptually quite similar to the suggestion of Harris (2013), it eliminates participants' incentives to send multiple messages because it allocates resources on a per-participant basis rather than a per-order basis. Thomson Reuters Matching chose this design rather than batch auctions so as to minimize the impact on participants whose existing trading strategies do not rely on speed.

Several of the formal models, including that of Du & Zhu (2017), build on a formulation introduced by Vives (2011) and Rostek & Weretka (2012) in which investors enjoy quasilinear utility and their payoff is equal to their investment return minus a holding cost that is a quadratic function of their portfolio holdings. The comparative statics of the earlier papers guide those of the later ones.

The modern market microstructure literature traces its roots to Kyle (1989) and Glosten & Milgrom (1985), whose papers can be read as responses to a famous puzzle about how traders' private information comes to be reflected in prices when traders are price takers. Following Grossman & Stiglitz (1980), suppose that traders spend resources to gather private information to earn trading profits. When traders are so informed, the equilibrium prices that clear the market depend on what the traders have learned, but that creates a tension with the assumption of price-taking behavior. In the extreme case where prices fully reflect all private information, no trader can earn any return on its investment in information, so any equilibrium must involve no such investment and prices that are unresponsive to information. However, that gives each trader an incentive to invest in information. The conclusion is that no equilibrium can exist with market clearing and endogenous information gathering in that extreme case. Grossman & Stiglitz resolve this puzzle by changing the economic environment: They add noise traders whose random behavior allows informed traders to earn trading profits. The market microstructure literature offers a different resolution. In the Glosten-Milgrom model (Glosten & Milgrom 1985), for example, although traders are price takers in the sense that they take the bid and ask prices as given, the prices at which they trade depend on whether they buy or sell and hence depend on the traders' private information.

CRYPTOCURRENCIES

Recent years have seen a huge upsurge of academic interest in and public awareness of cryptocurrencies. In the press, many have hailed cryptocurrencies and their associated public blockchain technology as innovations with the potential to change the nature of financial institutions. Among the recent applications being explored or developed are blockchain-based stock exchanges; fast, low-cost international bank transfers; smart contracts that can automatically execute contingent transactions; improved online identity management and security; and internal accounting within large companies. In each of these examples, advocates have argued that blockchain technology has the potential to improve speed and transparency and decrease costs by reducing the need for intermediaries, which could be especially beneficial to consumers in countries with underdeveloped financial systems. The advantages claimed for blockchain-based markets—increased speed, greater privacy, and lower transaction costs—are ones that expand markets but not ones that go far outside the bounds of traditional economic theorizing.

The first cryptocurrency, Bitcoin, and its associated blockchain technology were introduced by Satoshi Nakamoto (2008) and have attracted both imitators and increased academic interest over the past five years. Böhme et al. (2015) discusses Bitcoin's institutional details, its uses (including illegal activities), and the risks associated with using Bitcoin as a currency. Yermack (2015) summarizes the history of Bitcoin and argues that Bitcoin behaves less like a currency and more like a speculative investment. To resolve the mystery of what determines the price of Bitcoins, Athey et al. (2016) develop a theoretical model of cryptocurrency adoption for when users face exchange rate uncertainty. That model links Bitcoin exchange rates to market fundamentals, but the paper's empirical analysis yields mixed results and suggests, contrary to the theory, that the majority of Bitcoins are held by investors and infrequent users. Babaioff et al. (2016) analyzes the incentives for Bitcoin nodes to propagate information regarding transactions to neighboring nodes in the Bitcoin network. While such information propagation is critical for network function, the current protocol does not appropriately incentivize nodes. The authors propose an augmented protocol, which would better reward information propagation. Chiu & Koeppl (2017) treat Bitcoin mining as a rent-seeking contest in a monetary system model. Calibrating the model using Bitcoin data, the key finding is that the deadweight loss from rent seeking depends on how transaction fees are set under the Bitcoin protocol. They suggest that the loss could be reduced by using monetary growth to fund the system's infrastructure. Ciaian et al. (2016) study how the price of Bitcoin is determined by studying the empirical determinants of supply and demand. Gandal & Halaburda (2016) offer an empirical analysis of competition between cryptocurrencies using price data. They find that, after an introductory period has passed, the data are consistent with the use of cryptocurrencies as financial assets. Gans & Halaburda (2015) analyze private digital currencies issued by platforms such as Facebook, Microsoft, and Amazon to increase user activity as a means to increase advertising revenue.

Huberman et al. (2017) analyze Bitcoin as a payment system supported by a ledger maintained by a group of independent, profit-maximizing miners. Operating without the benefit of a trusted central coordinating authority, the integrity of the system is maintained by a combination of cryptography and suitable incentives. There are lags in the execution of user orders, and users bid for priority to get fast execution. Congestion is necessary to generate positive revenues.

Because the Bitcoin system is decentralized and does not rely on any trusted central authority, any disagreements about transactions must be resolved by a voting mechanism. Budish (2018) analyzes the security of the Bitcoin system from a so-called majority attack. The gain from such an attack grows with the size of the system, so to deter an attack, the cost of acquiring a majority share by deploying many miners must grow proportionately. It follows that the system cannot enjoy large-scale economies, which limits the potential of Bitcoin.

PROCUREMENT AUCTIONS

In most auction theory presented by economists, the analysis is focused on prices alone rather than on details of delivery or performance of the good or service. In those models, there is no significant difference between the auctions conducted by buyers and sellers. When a bidder with value v for an item bids and wins at price b, its payoff is $v - b$. When a bidder with cost c bids and wins a supply contract at price p, its payoff is $p - c$. With a simple sign change (Milgrom 2004), setting $c = -v$ and $p = -b$, the payoff formulas become identical and the auctioneer's payoff is similarly described by a simple sign change.

Issues of performance and default can arise whether the auctioneer is a buyer or a seller. Some examples of very large defaults by buyers include the Nextwave bankruptcy and default in the United States in 1996 on roughly \$10 billion of spectrum licenses won at auction; the \$19 billion auction sale of Telebras in 1998, for which the price was later renegotiated by the winning bidders; and the ITV Digital bankruptcy in 2002, after it had bid £178 million at auction to acquire soccer broadcast rights in the United Kingdom. Given the possibility of default, even an auctioneer who is a seller may wish to select a winning bidder by considering a score that depends both on the winning bid and on the bidder's identity or characteristics. There can be additional reasons to do that when the auctioneer is a buyer, since the quality and dependability of the good or service it acquires may depend on the bidder.

The analysis of scoring in procurements received its first serious attention from Che (1993), who studied a model in which a bid consisted of a price and a product (a point in characteristic space). The paper finds that, if the auctioneer fails to make the scoring of characteristics clear before the auction, bidders will be incentivized to differentiate their offerings excessively and raise their prices, hoping that they can win with a high price because the auctioneer happens to favor their design. The auctioneer does best when it can dash those hopes of a high profit by making clear how it will score the product characteristics of any seller's bid.

Bajari et al. (2008) study a buyer's choice between negotiating the purchase of a customized good or service versus specifying a standard good or service and using competitive bidding to find the lowest price. Theoretically, negotiations are most favored relative to auctions when projects are complex, which promotes highly incomplete contracts, or when the number of potential bidders is small. The paper tests these conclusions using a comprehensive data set of private sector building contracts awarded in Northern California during the years 1995–2000. The empirical findings are generally consistent with the theoretical account.

Arozamena & Cantillon (2004) investigate how firms' incentives to invest in observable cost reduction vary with the auction rule. In a second-price auction, a firm's investment does not affect the competitor's bid, so the competitor's response does not affect investment incentives. However, in a first-price auction, an investment that reduces a firm's costs may tend to make its competitor reduce its bid. In the comparison, investment incentives are weaker in the first-price auction than in the second-price auction.

Burguet et al. (2012) emphasize the importance of bankruptcy in procurement auctions. In the United States during 1990–1997, more than 80,000 contractors filed for bankruptcy, leaving unfinished construction projects. More than 60% of failure-to-complete construction projects resulted from such failures.

The paper presents a theory of procurement auctions when firms have limited liability and uncertain costs of performing the project. The main intuition is best illustrated by an extreme case in which firms have identical costs but differ in their financing. The best-financed firms lose more when they suffer an adverse cost shock, so they bid most cautiously, and the winning bidder is the firm with the poorest financial capacity. In this setting, selecting the lowest bidder leads to selecting the firm that is most likely to default. The same effect is still at work even if it may be muted by variations in the expected cost of performance. It is still not possible, using an auction, to select the firm with the lowest cost as the winner.

Board (2007) investigates the interaction between the auction rules (first-price or second-price) and limited liability. This paper finds that a second-price auction induces higher prices, higher bankruptcy rates, and lower bidder utility than a first-price auction. The paper also shows that the handling of bankruptcy critically affects the analysis. It considers cases in which bankruptcy can lead to large or small recovery or resale, in which the assets of the contracting firm remain intact and can be sold to losers of the original auction. If bankruptcy is sufficiently costly, the auctioneer prefers a first-price, while if bankrupt firms remain intact, that preference is reversed.

Some government procurements suggest the possibility that auction design decisions have been captured by the regulated bidders. Cramton et al. (2015) study the design of the auction commonly used in the United States for Medicare durable medical equipment. Two unusual features of these auctions are that (*a*) winners are paid the median winning bid and (*b*) bids are nonbinding. The authors show that this auction leads to inefficiencies in two ways. First, in equilibrium, some firms may refuse to supply, as the price is set below their cost. Second, some high-cost firms potentially displace low-cost firms as winners.

COMBINATORIAL AUCTIONS

Whether the auctioneer is a buyer or seller, one very common issue is that the sale may involve many complementary purchases or sales, often arising from fixed costs or other scale economies. For example, a supplier shipping products to a location may enjoy economies of scale by shipping in a container-load and may offer a discount to a customer who buys a full container of items. Economies of scale are also relevant factors in many of the markets discussed earlier in this review. For example, in electricity markets, the cost of power from a generation facility involves fixed costs, so the operator is willing to supply at a lower price if more power is being purchased. Similarly, a mobile phone operator that is purchasing radio spectrum rights might find it more valuable to acquire rights in the same frequency band in every area that it serves because of scale economies that can be achieved in developing and setting up its complete network to use that frequency band.

With nonconvexities and scale economies, it comes as little surprise that market-clearing prices may fail to exist. That means that there may be no set of prices for individual items that clear a particular market, but there may still be pricing based on combinations of items that could work. As discussed earlier, many recent procurement auctions and radio spectrum auctions are combinatorial auctions.

One combinatorial auction design that is often used as a reference point is the Vickrey auction. It has the property that so long as bidders (*a*) care only about what they acquire, (*b*) know their own values for packages of items, (*c*) have quasilinear payoffs, and (*d*) are not effectively budget-constrained, they have no incentive to misrepresent their values to the auctioneer. Moreover, when they report truthfully, the auction selects the efficient outcome. If the auction entails no payments to losers, then the theorems of Green & Laffont (1977) and Holmstrom (1979) establish that, even on limited (but connected) preference domains, the Vickrey auction is the only mechanism with these properties. These properties often make the Vickrey design the standard against which other designs are tested.

Still, the Vickrey design has important drawbacks. One is that the Vickrey allocation and pricing may result in payoffs that are not in the core of the game involving the bidders and the seller. The reason is, roughly, that Vickrey prices can be uncompetitively low. Milgrom (2007) illustrates this with a simple and extreme example: Suppose that a seller has two units for sale. There are two bidders who are each willing to pay 1 for either one or two units of the good and a third bidder who will pay nothing for a single unit but is willing to pay 1 for the two units together. In this example, even though there are three bidders, each of whom would be willing to pay 1 for the two units, the Vickrey prices are zero! Ausubel & Milgrom (2006) elaborate on this further, showing that the arithmetic of Vickrey pricing makes the auction unusually vulnerable to collusion, shill bidding, and other ills. Repeating a result from their previous work (Ausubel & Milgrom 2002), they also show that if the goods are substitutes for all the bidders, then the Vickrey outcome is in the core; that is, Vickrey prices are competitive. Another important problem for a Vickrey auction for N distinct items is the challenge it poses to bidders who must determine and communicate values for all $2^N - 1$ combinations of items, especially when $N \geq 10$, and the complexity of the computation of Vickrey prices increases substantially when the number of items is much larger.

The book *Combinatorial Auctions*, edited by Cramton et al. (2006), includes the Ausubel & Milgrom paper and others. While research in combinatorial auctions has continued among computer scientists concerned about the complexity of the computations and communications required by the Vickrey auction, many of the main economic papers were published before 2006. The package assignment model of Bikhchandani & Ostroy (2002) formulates the general problem and characterizes what kinds of linear or nonlinear prices may be needed to support the efficient solution. Sometimes even nonlinear prices for packages of goods cannot do the job unless different individuals face different prices, suggesting that efficient allocations cannot be decentralized.

De Vries & Vohra (2003) provide a survey of combinatorial auctions emphasizing mathematical and algorithmic issues. They highlight three such issues: (*a*) the problem of finding a bidding language that is suitable for taming the potentially extreme complexity of combinatorial bidding; (*b*) the winner determination problem, which can be a large-scale integer programming problem that is hard to solve; and (*c*) the pricing problem, that is, assigning values to individual items when no prices exist to support the optimal allocation.

Several papers treat applications of combinatorial auctions in a variety of settings. Caplice & Sheffi (2006) study the application of sealed-bid package auctions to allocate bus routes in London, England. Epstein et al. (2002) study the use of package bidding to acquire milk supplies for schools in Chile, including schools in remote areas. Bichler et al. (2006) study the use of combinatorial auctions for industrial procurement, including packaging, chemicals, and road construction and repair. They report that the important issues in industrial combinatorial auctions are related to representative bidding language, winner determination, and side constraints on the allocation. Hohner et al. (2003) introduce a bidding website for combinatorial package auctions in procurement applications designed by Mars-IBM. The auction website allows for a variety of bid structures, including bundled all-or-nothing bids and quantity-discounted bids. Users of this website can also submit multiattribute bids.

CONCLUSION

As a human activity, market design is an old topic. Even ancient societies had to agree on the days and places for markets so that people knew where and when to show up, and they needed to set some rules to govern trade. As interregional trade grew in Europe, the Champagne fairs developed commercial laws, traditions, and means to enforce contracts and to meet and identify reliable trading partners. In the modern era, with electronic, online trading, an early task was to update these historical activities. Today's meeting places are online exchanges, including financial exchanges, matching apps for workers, ride-sharing, dating, and more. As fast, low-cost trading became possible, adaptations became necessary, such as rules to combat the evils of high-frequency trading. New kinds of exchanges and markets to support them also emerged. Combinatorial auctions have become routine for procurement, electricity, radio spectrum, etc., and along with these, new auction rules have evolved for setting package prices. Rules have evolved to make it easier for nonhuman agents to bid effectively in auctions, as needed for Internet advertising.

Recent years have witnessed many changes taking advantage of new technologies to expand the role of markets. New theories have begun to emerge, which do not conform to the old assumptions of neoclassical economics, to understand the challenges facing these markets. Not every good in a category is identical, and part of the progress of market design has been splitting categories to allow much finer matching of items to buyers (as in Internet advertising) and much more accurate pricing (e.g., electricity according to the time of day). New theories explain how package pricing can be used to support more efficient resource allocation when goods are complements or scale economies are present; how cryptocurrencies survive or fail; how low-cost, high-speed trading disrupts traditional markets; and more. With these new theories comes understanding that may help us to design more effective markets and avoid or regulate the problems that arise. All in all, it is an exciting time for market design.

DISCLOSURE STATEMENT

The author is not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

The author thanks Matt Jackson and his research assistants, Negar Matoorian Pour, Ellen Muir, and Zi Yang Kang, for detailed comments.

LITERATURE CITED

Amaldoss W, Jerath K, Sayedi A. 2016. Keyword management costs and "broad match" in sponsored search advertising. *Mark. Sci.* 35:259–74

Arnosti N, Beck M, Milgrom P. 2016. Adverse selection and auction design for Internet display advertising. *Am. Econ. Rev.* 106:2852–66

Arozamena L, Cantillon E. 2004. Investment incentives in procurement auctions. *Rev. Econ. Stud.* 71:1–18

Arrow KJ, Block HD, Hurwicz L. 1959. On the stability of the competitive equilibrium, II. *Econometrica* 27(1):82–109

Athey S, Ellison G. 2011. Position auctions with consumer search. *Q. J. Econ.* 126:1213–70

Athey S, Parashkevov I, Sarukkai V, Xia J. 2016. *Bitcoin pricing, adoption, and usage: theory and evidence*. Work. Pap. 3469, Stanford Univ. Bus. School, Stanford, CA

Ausubel LM, Cramton P, Milgrom P. 2006. The clock-proxy auction: a practical combinatorial auction design. See Cramton et al. 2006, pp. 115–38

Ausubel LM, Cramton P, Pycia M, Rostek M, Weretka M. 2014. Demand reduction and inefficiency in multi-unit auctions. *Rev. Econ. Stud.* 81:1366–400

Ausubel LM, Milgrom P. 2002. Ascending auctions with package bidding. *Adv. Theor. Econ.* 1(1). **https://doi.org/10.2202/1534-5963.1019**

Ausubel LM, Milgrom P. 2006. The lovely but lonely Vickrey auction. See Cramton et al. 2006, pp. 22–26

Babaioff M, Dobzinski S, Oren S, Zohar A. 2016. On Bitcoin and red balloons. arXiv:1111.2626 [cs.GT]

Bajari P, McMillan R, Tadelis S. 2008. Auctions versus negotiations in procurement: an empirical analysis. *J. Law Econ. Organ.* 25:372–99

Biais B, Foucault T, Moinas S. 2015. Equilibrium fast trading. *J. Financ. Econ.* 116:292–313

Bichler M, Davenport A, Hohner G, Kalagnanam J. 2006. Industrial procurement auctions. See Cramton et al. 2006, pp. 593–612

Bichler M, Goeree JK, eds. 2017. *Handbook of Spectrum Auction Design*. Cambridge, UK: Cambridge Univ. Press

Bikhchandani S, Ostroy JM. 2002. The package assignment model. *J. Econ. Theory* 107:377–406

Board S. 2007. Bidding into the red: a model of post-auction bankruptcy. *J. Finance* 62:2695–723

Böhme R, Christin N, Edelman B, Moore T. 2015. Bitcoin: economics, technology, and governance. *J. Econ. Perspect.* 29:213–38

Brusco S, Jackson MO. 1999. The optimal design of a market. *J. Econ. Theory* 88(1):1–39

Budish E. 2018. *The economic limits of Bitcoin and the blockchain*. Work. Pap., Univ. Chicago Booth Bus. School, Chicago, IL

Budish E, Cramton P, Shim J. 2015. The high-frequency trading arms race: frequent batch auctions as a market design response. *Q. J. Econ.* 130:1547–621

Budish E, Lee R, Shim J. 2018. *Will the market fix the market? A theory of stock exchange competition and innovation*. Work. Pap., Univ. Chicago Booth Bus. School, Chicago, IL

Burguet R, Ganuza JJ, Hauk E. 2012. Limited liability and mechanism design in procurement. *Games Econ. Behav.* 76:15–25

Caplice C, Sheffi Y. 2006. Combinatorial auctions for truckload transportation. See Cramton et al. 2006, pp. 539–71

Charles River Associates, Market Design, Inc. 1997. *Package bidding for spectrum licenses*. Rep., US Fed. Commun. Comm., Washington, DC

Che Y-K. 1993. Design competition through multidimensional auctions. *RAND J. Econ.* 24(4):668–80

Chiu J, Koeppl T. 2017. *The economics of cryptocurrencies—Bitcoin and beyond*. Work. Pap. 1389, Dep. Econ., Queen's Univ., Kingston, Ontario

Ciaian P, Rajcaniova M, Kancs d'A. 2016. The economics of Bitcoin price formation. *Appl. Econ.* 48:1799–815

Coase RH. 1959. The Federal Communications Commission. *J. Law Econ.* 2:1–40

Cramton P. 2013. Spectrum auction design. *Rev. Ind. Organ.* 42:161–90

Cramton P. 2017. Electricity market design. *Oxf. Rev. Econ. Policy* 33:589–612

Cramton P, Ellermeyer S, Katzman B. 2015. Designed to fail: the Medicare auction for durable medical equipment. *Econ. Inq.* 53:469–85

Cramton P, Schwartz JA. 2002. Collusive bidding in the FCC spectrum auctions. *Contrib. Econ. Anal. Policy* 1:11

Cramton P, Shoham Y, Steinberg R, eds. 2006. *Combinatorial Auctions*. Cambridge, MA: MIT Press

Day RW, Cramton P. 2012. Quadratic core-selecting payment rules for combinatorial auctions. *Oper. Res.* 60:588–603

Day RW, Milgrom P. 2008. Core-selecting package auctions. *Int. J. Game Theory* 36:393–407

Day RW, Raghavan S. 2007. Fair payments for efficient allocations in public sector combinatorial auctions. *Manag. Sci.* 53:1389–406

De Vries S, Vohra RV. 2003. Combinatorial auctions: a survey. *INFORMS J. Comput.* 15:284–309

Dhangwatnotai P. 2011. Multi-keyword sponsored search. In *Proceedings of the 12th ACM Conference on Electronic Commerce, San Jose, CA, June 5–9, 2011*, pp. 91–100. New York: ACM

Du S, Zhu H. 2017. What is the optimal trading frequency in financial markets? *Rev. Econ. Stud.* 84:1606–51

Edelman B, Ostrovsky M. 2007. Strategic bidder behavior in sponsored search auctions. *Decis. Support Syst.* 43:192–98

Edelman B, Ostrovsky M, Schwarz M. 2007. Internet advertising and the generalized second-price auction: selling billions of dollars worth of keywords. *Am. Econ. Rev.* 97:242–59

Eliaz K, Spiegler R. 2016. Search design and broad matching. *Am. Econ. Rev.* 106:563–86

Epstein R, Henríquez L, Catalán J, Weintraub GY, Martínez C. 2002. A combinational auction improves school meals in Chile. *Interfaces* 32:1–14

Evans DS. 2009. The online advertising industry: economics, evolution, and privacy. *J. Econ. Perspect.* 23:37–60

Fabra N, von der Fehr N-H, Harbord D. 2006. Designing electricity auctions. *RAND J. Econ.* 37:23–46

Gandal N, Halaburda H. 2016. Can we predict the winner in a market with network effects? Competition in cryptocurrency market. *Games* 7(3):16. **https://doi.org/10.3390/g7030016**

Gans JS, Halaburda H. 2015. Some economics of private digital currency. In *Economic Analysis of the Digital Economy*, ed. A Goldfarb, SM Greenstein, CE Tucker, pp. 257–76. Chicago, IL: Univ. Chicago Press

Glosten LR, Milgrom PR. 1985. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *J. Financial Econ.* 14:71–100

Green J, Laffont J-J. 1977. Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica* 45(2):427–38

Grossman SJ, Stiglitz JE. 1980. On the impossibility of informationally efficient markets. *Am. Econ. Rev.* 70:393–408

Harris L. 2013. What to do about high-frequency trading. *Financial Anal. J.* 69:6–9

Herzel L. 1951. "Public interest" and the market in color television regulation. *Univ. Chicago Law Rev.* 18:802–16

Hohner G, Rich J, Ng E, Reid G, Davenport AJ, et al. 2003. Combinatorial and quantity-discount procurement auctions benefit Mars, Incorporated and its suppliers. *Interfaces* 33:23–35

Holmstrom B. 1979. Moral hazard and observability. *Bell J. Econ.* 10:74–91

Hu Y, Shin J, Tang Z. 2015. Incentive problems in performance-based online advertising pricing: cost per click versus cost per action. *Manag. Sci.* 62:2022–38

Huberman G, Leshno JD, Moallemi CC. 2017. *Monopoly without a monopolist: an economic analysis of the Bitcoin payment system*. Work. Pap. 27/2017, Bank of Finland, Helsinki

Joskow PL, Wolfram CD. 2012. Dynamic pricing of electricity. *Am. Econ. Rev.* 102:381–85

Kelso AS Jr., Crawford VP. 1982. Job matching, coalition formation, and gross substitutes. *Econometrica* 50:1483–504

Klemperer P. 2002. What really matters in auction design. *J. Econ. Perspect.* 16:169–89

Kyle AS. 1989. Informed speculation with imperfect competition. *Rev. Econ. Stud.* 56:317–55

Kyle AS, Lee J. 2017. Towards a fully continuous exchange. *Oxf. Rev. Econ. Policy* 33:650–75

Lahaie S, Pennock DM, Saberi A, Vohra RV. 2007. Sponsored search auctions. In *Algorithmic Game Theory*, ed. N Nisan, T Roughgarden, E Tardos, V Vazirani, pp. 699–716. Cambridge, UK: Cambridge Univ. Press

Ledyard JO, Porter D, Rangel J. 1994. Using computerized exchange systems to solve an allocation problem in project management. *J. Organ. Comput.* 4(3):271–96

Levin J, Milgrom P. 2010. Online advertising: heterogeneity and conflation in market design. *Am. Econ. Rev.* 100:603–7

Lewis M. 2014. *Flash Boys: A Wall Street Revolt*. New York: W.W. Norton

Leyton-Brown K, Milgrom P, Segal I. 2017. Economics and computer science of a radio spectrum reallocation. *PNAS* 114:7202–9

Li S. 2017. Obviously strategy-proof mechanisms. *Am. Econ. Rev.* 107(11):3257–87

Maurer L, Barroso L. 2011. *Electricity auctions: an overview of efficient practices*. Rep. 63875, World Bank, Washington, DC

McAdams D. 2018. *Smart watershed markets: the case of the Central Platte groundwater exchange in Nebraska*. Work. Pap., Econ. Dep., Duke Univ., Durham, NC

McMillan J. 1994. Selling spectrum rights. *J. Econ. Perspect.* 8:145–62

Melton H. 2017. Market mechanism refinement on a continuous limit order book venue: a case study. *ACM SIGecom Exchanges* 16:74–79

Milgrom P. 2000. Putting auction theory to work: the simultaneous ascending auction. *J. Political Econ.* 108:245–72

Milgrom P. 2004. *Putting Auction Theory to Work*. Cambridge, UK: Cambridge Univ. Press

Milgrom P. 2007. Package auctions and exchanges. *Econometrica* 75:935–65

Milgrom P. 2017. *Discovering Prices: Auction Design in Markets with Complex Constraints*. New York: Columbia Univ. Press

Milgrom P, Ausubel L, Levin J, Segal I. 2012. *Incentive auction rules option and discussion*. Rep., US Fed. Commun. Comm., Washington, DC

Milgrom P, Mollner J. 2018a. Equilibrium selection in auctions and high stakes games. *Econometrica* 86:219–61

Milgrom P, Mollner J. 2018b. *Extended proper equilibrium*. Work. Pap., Stanford Univ., Stanford, CA

Milgrom P, Segal I. 2018. Clock auctions and radio spectrum reallocation. *J. Political Econ.* In press

Nakamoto S. 2008. *Bitcoin: a peer-to-peer electronic cash system*. White Pap., Bitcoin. **https://bitcoin.org/bitcoin.pdf**

PricewaterhouseCoopers. 2018. *IAB Internet advertising revenue report*. Rep., Interact. Advert. Bureau, New York. **https://www.iab.com/wp-content/uploads/2018/05/IAB-2017-Full-Year-Internet-Advertising-Revenue-Report.REV_.pdf**

Rostek M, Weretka M. 2012. Price inference in small markets. *Econometrica* 80:687–711

Rostek M, Weretka M. 2015. Dynamic thin markets. *Rev. Financial Stud.* 28:2946–92

Roth AE. 2016. *Who Gets What—and Why: The New Economics of Matchmaking and Market Design*. New York: Mariner

Sayedi A. 2018. Real-time bidding in online display advertising. *Mark. Sci.* 37:553–68

Smith VL. 1962. An experimental study of competitive market behavior. *J. Political Econ.* 70:111–37

Smith VL. 1965. Experimental auction markets and the Walrasian hypothesis. *J. Political Econ.* 73(4):387–93

Varian HR. 2007. Positional auctions. *Int. J. Ind. Organ.* 25:1163–78

Vayanos D. 1999. Strategic trading and welfare in a dynamic market. *Rev. Econ. Stud.* 66:219–54

Vives X. 2011. Strategic supply function competition with private information. *Econometrica* 79:1919–66

Wang M. 2018. *Header bidding and auctions with intermediaries*. Undergrad. Honors Thesis, Stanford Univ., Stanford, CA

Wilson R. 2002. Architecture of power markets. *Econometrica* 70:1299–340

Yermack D. 2015. Is Bitcoin a real currency? An economic appraisal. In *Handbook of Digital Currency*, ed. DLK Chuen, pp. 31–43. London: Elsevier/Academic Press