Auction Market Design: Recent Innovations

Paul Milgrom

August 2018

Introduction

Introductory economics courses begin to teach about markets by focusing on trades of a single homogeneous good in which 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 require market design effort to deal with cases in which the identities of trading partners matter; prices play little or no role in market clearing; or prices vary greatly among similar transactions. There are many examples. In the marriage market, men and women care deeply about the identity of their dating or marriage partners, even when there are no dowries or bride-prices involved. Markets for matching children to schools or kidney donors to kidney patients make no use of money or prices to guide the allocation. Markets for advertising on Internet search pages use auctions to determine a separate price for each ad impression.

A prominent feature of many markets is that fine distinctions are made among similar-seeming goods. The most relevant ad to show an Internet user varies not only user-by-user, but even over short periods of time for a single user. The most relevant ad to show to a user varies over time, depending on whether the user is currently shopping to replace a broken kitchen appliance, refinance a mortgage, or plan a vacation. In many electrical power markets, power is distinguished by the time and location at which energy is to be made available, with shorter time intervals and smaller geographic areas becoming increasingly common.

Another way that some real markets differ from textbook treatments is that they may solve economic problems that include more constraints than just the ones that limit the quantity of each good available. Consider, for example, the problem of assigning television stations to broadcast channels to be used for over-the-air broadcasting. We can represent this problem graphically by treating each TV station as a node in the graph and connecting two stations by an arc if those stations cannot be assigned to the same channel without creating interference. Any feasible allocation assigns channels to stations so that no two connected stations 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 they cannot be efficiently solved by finding prices to assign to each constraint. That suggests that organizing a market to use prices to guide these transactions is unlikely to be lead to a good solution.

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 and Hurwicz, 1959)** that prices which are adjusted separately in each market can promote simultaneous clearing in all markets, but complements are more challenging. Tightly connected complements are most challenging of all. For example, a 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, and 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 some 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 its 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 enough.

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 to 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 also easier 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. Given the recency of the developments, understandings are evolving rapidly and many important papers are still unpublished.

This review seeks to highlight some of the questions and challenges found in modern literature and modern marketplaces and how those have been addressed in theory and practice. The paper is organized around the different kinds of applications and how market rules developed to solve real resource allocation problems.

Internet Advertising

Before the year 2000, the Internet was mostly regarded as just one more advertising medium, similar to print or radio or television: Internet publishers sold to advertisers by contract with prices based largely on the number of impressions the site could deliver. 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% (**PwC, 2018**).

The first big changes to distinguish advertising on the Internet came with *sponsored search*, according to which an advertiser could pay to have its ads appear alongside organic search results when a user searched for particular *keywords*. For example, an auto insurance company might select the "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 would 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. **[Which citation for this history?]**

Sponsored searches quickly began to generate huge revenues for Google, and that inspired other Internet publishers to search for 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 to use 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 hotel, restaurants and attractions. Publishers and others developed systematic ways to report and share information, so a user who bought a ticket on an airline website could see the travel-related ad on a completely unrelated website, such as ones for sporting news or home furnishings. 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 that had not been a problem for traditional advertising. 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? Second, 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 [**Which citation??**]. 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 an economic problem. The solution is not a perfect one: some clicks may be fraudulent ones by dishonest publishers or by 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 predicting the click rates of different ads.

Another innovation that search engines developed was *broad match* **[Which citation??]**, which is best described by an example. Suppose a company specifies "property insurance" as a keyword while the user searches for something else, like "fire insurance" or "protection against property damage." Broad match allows the search site to decide which of these related terms should trigger the ad. If the advertiser gets value from clicks, then it could be willing to delegate such a decision. And, the search engine has good incentives to predict the click rate of these broad matches, in order to show the right ads to maximize its revenue.

The next topic is auction rules for sponsored search. In the early history of sponsored search auctions, winners paid the prices that they had bid. To see why that became problematic, suppose that there are two bidders for two positions on a search page. Imagine that one bids $2/thousand impressions, a second bids $1, and the reserve price is also $1. The $2-bidder wins the top position and the $1-bidder wins the second position. So far, so good. However, the $2-bidder may soon learn that it can still win the top position even if it reduces its bid to $1.01. Once it makes that reduction, the second bidder may learn that it can raise its bid from $1 to $1.02 to win the first position instead of the second. The first bidder might then raise its bid to $1.03, and the price continues to climbs until one of the bidders, perhaps one who has just paid $1.99 to win the second position, realizes that rather than raising its bid to $2.01 to win the first position, it can *reduce* its bid to $1 and still win second position, getting a better deal for itself. If it does that, the whole cycle starts over again. Since this is happening on the Internet, the bid revision processes were often automated and the cycles were very visible in data, which was compiled and graphed by **Edelman and Ostrovsky (2006)**.

This instability of prices and winners can be a problem for many reasons, including the inefficient allocation, in which 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 bid $2 and another bid $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 wins 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 was studied by **Edelman, Ostrovsky and Schwartz (2007?)** and **Varian (2007?)** using non-cooperative 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. **Milgrom and Mollner (2018a, 2018b)** criticize these analyses as *ad hoc* for failing to apply general game theoretic reasoning to select an equilibrium. They show that these same conclusions about prices can be derived using their *test-set* and *extended proper* equilibrium refinements, which apply to general finite games, but only if the parameters of the auction setting satisfy particular restrictions.

Athey and Ellison (2011) emphasize that besides the advertiser and the publisher, there is another important party in an ad auction: the user. Relevant ads improve consumer experience, helping them to find the sellers from whom they prefer to buy. **Eliaz and Spiegler (2016)** show how a properly design broad match algorithm can improve consumer search and, under special conditions, can be an efficient indirect mechanism for extracting value for the publisher. Generally, broad match (like Smart Pricing, as discussed below) is a way to *simplify* bidding, allowing advertisers to express values simply for one narrow set of outcomes while the market mechanism infer values for other outcomes and promotes those. The market design literature has still paid only limited attention to simplification, with the main general exception being **Levin and Milgrom (2010)**, which studies how “conflation” - the suppressing of distinctions - can sometimes simplify and lubricate the operation of markets.

The generalized second-price auction applies to sponsored search advertising, which currently accounts for about 45% of all Internet advertising revenue, a volume that is now rivaled by Internet display advertising. In the past, display ads on the Internet were sold either through contracts or by second-price auctions, but advertisers have limited the sites they bid on because of concerns about quality. There is a huge variety in the character of sites on the Internet, which include pornography sites, financial news, barbecue recipes, and much else, attracting very different categories of users and engaging them in very different ways. To be informed participants in automated ad auctions, advertisers need to know about each site and user, distinguishing among impressions and clicks to estimate their values and determining their per-click bids from that assessment. Lacking good information to place accurate bids on the full range of websites, advertisers were initially cautious, often limiting themselves just to search advertising or to search plus a few carefully selected websites. Such behavior results in thinner markets, poorer matches of advertisers to users, and lower auction prices and calls for improved market design.

One early attempt to resolve this information problem was Google’s *Smart Pricing* program, which was introduced in 2004 (**citation?**). Each participating advertiser in the Smart Pricing program would specify its own definition of *good performance* of a click. For example, good performance might mean that the click resulted in a sale or in the user completing a form or visiting a second page on the advertiser's website. Operating under a contract, the advertiser would allow Google to monitor the performance of clicks on its website in order to predict the performance of future clicks on various websites and to compare them to the performance of a sponsored search click. Smart Pricing would start with the advertiser's sponsored-search bid per click and adjust it for each website based on the expected performance of clicks from that site, greatly reducing the effort that an advertiser must expend to expand the reach of its bidding beyond sponsored search.

Smart Pricing can be regarded as a response to a traditional challenge for markets: the problem of *adverse selection* which arises when a class of market participants, whom we may call "performance advertisers," cannot predict the performance of their ad impressions. **Arnosti, Beck and Milgrom (2016)** investigate adverse selection against *contract advertisers* who suffer an information disadvantage compared to performance advertisers, because they are not looking for performance on the Internet. For example, a shopping mall may effectively advertise its new weekend hours without inspiring any user to click on its ad, but it cannot tell whether consumers show up because of this ad or because of other promotion efforts.

To capture these ideas, the formal model includes two kinds of advertisers. One is a contract advertiser to whom a publisher has committed to sell some fixed number (or fraction) of its relevant impressions. This contract advertiser can anticipate the average value of a random impression in that set but cannot observe what it gets. There are also two or more *performance advertisers*, each of whom knows its own value for any individual impression. The value of impression i to advertiser j is the product of two independent random variables: $x_i v_j$. In this formulation, the x_i variable represents value that depends on attributes of the user, such as the user's income and propensity to respond to Internet display ads. A higher value of x_i makes the impression more valuable to everyone, including the contract advertiser. The v_j variable is different for different advertisers, reflecting the different quality of different matches. In this formulation, the second-price auction traditionally used for allocating Internet display ad impressions leads to good matching for performance advertisers, but also to adverse selection against the contract advertiser, who would expect to lose most of the high x_i impressions.

The paper studies whether there is an alternative design that might have better properties. It approaches the question axiomatically. The goal is to find a design that has these properties:

- It is *adverse-selection free*. This means that the distribution of x_i values in the selection of impressions won by the contract advertiser should be identical to the distribution of x_i values in the full population.
- It is *efficient among performance advertisers*. That is, if a performance advertiser wins, then the winner is the one with the highest match value.
- It is *strategy-proof*. That is, it is a dominant strategy for performance bidders to set their bids equal to their actual values.
- It is *anonymous*. This means that the rule is symmetric and applies to any number of performance advertisers.
- It is *false-name proof*. No bidder should ever be able to reduce its price by submitting an extra bid (for example, at a low price) and the seller should never be able to raise the price by submitting an extra bid, provided that bid is no higher than the second highest bid.

The paper shows that there is a unique family of auctions, which they call the "modified second-bid auctions," that has all these 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 award to the highest bidder at a price equal to α times the second highest bid. Otherwise, the impression is awarded to the contract bidder.

The paper also includes a numerical analysis of the performance of the modified second-bid auction. It assumes that the match values are independently and identically distributed and drawn from a power law distribution, so that most of the value of good matching comes from very good matches. They assume that there is some fixed fraction of impressions that must be assigned to the performance advertiser. Their finding is that over all power law distributions, all distributions of x_i, and all fractions to be assigned to the contract advertiser, the worst-case ratio of the value achieved by the mechanism to the value of the full-information optimal matching exceeds 94%.

Several recent articles examine the mix of contract advertising and bidding from a different perspective. When bidders target just a few ads, there may be many opportunities with just one bid, leading to low revenue for publishers, but these bidders may offer to pay quite a lot for their advertisements. **Sayedi (2018)** study the optimal fraction of contract ads. In terms of the usual literature on auctions, contract ads serve a role similar to reserve prices. The right fraction of contract ads allows the seller to ensure that carefully targeted advertising, which delivers high value to performance ad buyers, also maximizes revenues for the seller.

The Internet advertising market is huge and quite dynamic, and new developments call for new solutions and new analytical research, beyond what is found in the current literature. One such development is variable size ads.

In the early days of sponsored search, all ads were restricted to be the same size, with a few lines of carefully chosen prose. Today, however, sponsored search ads can vary in size, from perhaps 3 to 18 lines. Given limited total number of ad lines on a page, publishers will need to charge more for larger ads and less for smaller ones. How should that be done? How can advertisers express their bids? How should winners be selected, and prices be determined?

Some initial attempts to deal with this have 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 hard problem (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 and 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 it 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, some publishers are soliciting bids from various ad exchanges and comparing those. This is a deeply problematic market organization. For example, suppose that the bids in exchange #1 are 5 and 3 while those in exchange #2 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 bids 5 beats out the bidder who bid 10. This organization creates an inefficiency and there is evidence that exchanges are changing their rules, shifting from second-price auctions to first-price auctions in order to attract bidders with higher values to their exchanges.

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 (1950)**, **Ronald Coase (1959)** became the most famous early advocate of using auctions instead, but his recommendations long fell on deaf ears. Critics reportedly laughed 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. 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 in **McMillan (1994)**. Someone needed to attend to the details!

In 1993, when the Congress first authorized the use of auctions for radio spectrum in the United States, the Federal Communications Commission (FCC), having no units with auction knowledge or experience, 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 issues a proposal that was based, to the extent possible, on academic citations. 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, bidding on a larger volume of licenses than in the previous round. Adapting the logic of **Kelso and Crawford (1982)** to this application, **Milgrom (2000)** showed 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 ensure 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 setting 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, Ausubel, Levin and Segal, 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)** argued that “what really matters in auction design” is traditional industrial policy to prevent with “collusive, predatory and entry-deterring behavior.” The importance of these elements is also affirmed by other analysts. **Cramton and Schwartz (2002)** find that bidders sometimes collude by using bids to signal, that is, to make threats and promises and to offer deals. **Ausubel, Cramton, Pycia, Rostek and Waretka, (2014)** find that bidders in an SMRA will often find it profitable to exercise market power, reducing demand to reduce the prices that they pay.

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, 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 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. That is a bidder's *exposure problem*, and laboratory experiments (Ledyard, Noussair and Porter, 1994) had begun to suggest that auction designs that avoided that problem might perform better 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 edited by **Cramton, Shoham and Steinberg (2006).** An article in that book by Ausubel, **Cramton and Milgrom (2006)** introduced a new auction format, now known as the "combinatorial clock auction" (CCA), which was designed to eliminate the exposure problem and mitigate the collusion problem.

The CCA works in two main stages. In its *clock* stage, it works much like the clock auction described 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, the offer a bidder makes is just to buy the whole package at the specified price. When the clock stage 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 combination of packages is the set that maximizes the total bid, but the prices are determined by a core-selecting pricing formula, according to the principle suggested by **Day and Raghavan (2007)**, **Day and Milgrom (2008)** and **Cramton and Day (2012)**. The resulting prices are typically close or identical to Vickrey prices and, in particular, each bidders' bids have little or no impact on the prices it pays for what it wins. In practice, this appears to eliminate any incentive for demand reduction, making it much harder for bidders to divide 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 US, that is easier to do when there is an 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 **(Milgrom and Segal, 2017)**. Optimizing the assignments of channels to stations is a large-scale, NP-complete problem that was beyond the capacity of even modern computers and software. This made it impractical to buy spectrum using a Vickrey auction, which must solve a separate optimization for each seller to determine that seller’s price. Also, Vickrey auctions seek efficiency without regard to cost, but cost control was a significant goal of the auction design.

The incentive auction used a new kind of descending clock auction to buy TV broadcast rights - one that offered a different price to each station depending on its characteristics. Using estimated stations values and simulating the auction results, **Leyton-Brown, Milgrom and Segal (2017)** compared the results of the auction to that of a Vickrey auction that seeks to minimize the total value of the stations removed from broadcasting. They found that the auction allocation, on average, adds about 5% to the minimum value of stations removed and reduces the cost of the procurement by about 24%.

Procurement Auctions

***

Electricity and Commodity Auctions

Bankruptcy Auctions

Crytocurrencies