Auction Research Evolving: Theorems and Market Designs*

Prize Lecture, December 7, 2020 by
Paul R. Milgrom
Stanford University, Stanford, CA, USA.

Game-theoretic modeling of auctions began in the 1960s with a pair of seminal papers by William Vickrey (1961, 1962) and the brilliant but unpublished doctoral dissertation of Armando Ortega-Reichert (1968). Robert Wilson (1977, 1979) became the next important contributor to auction theory research and, as Wilson's student, I was inspired to make auctions and bidding the subject of my doctoral dissertation.

My research about auction theory and market design has evolved through three distinct eras. In the first, which began with my dissertation and continued for about five years, I aimed to extend Wilson's research to a wider set of situations, characterizing the equilibrium strategies of auction games, the extent to which bidders' private information becomes reflected in prices in auctions and securities markets, and how a seller's expected revenue depends on the detailed auction rules. While the first era was about understanding existing rules, the next era was about designing new auctions. That era was launched when the FCC decided

* This paper expands my lecture delivered in recorded form during Nobel Prize Week 2020. A recording of that lecture is currently available on YouTube. I thank Andrew Vogt, Martin Bichler and Gregor Schwarz for their comments and suggestions.

to sell radio spectrum rights using an auction using the simultaneous multiple round (SMR) that Wilson and I had proposed. The FCC auction was designed in mere months and inspired me to study the properties of our SMR design and to explore alternatives. Nearly all the theory before the SMR era was about designing auctions for the sale of single items or items that were substitutes. In the third, current era, I designed practical auctions to cope with the new problems that arise when some items are not substitutes. The design and implementation teams I led used both theoretical and computational methods to achieve our aims, drawing on insights from economics, game theory and computer science.

I published mostly theoretical work in academic journals, proving theorems about the properties of mathematical models, but developing and participating in real-world mechanisms requires more than that. Two important lessons that I learned from working on high-stakes auctions are that they operate in an almost infinite variety of contexts and that this variety is the reason for the paradoxical importance of including unrealistic assumptions. No single set of assumptions is adequate to describe all the real settings that auction designers encounter, and too much specificity in models can blind the analyst to important general insights.

As one who uses theory both to prove theorems and to create mechanisms for high-stakes applications, I was initially puzzled by the fact that the most celebrated theorems of economic theory rely on deeply unrealistic assumptions. Arrow's and Debreu's first Welfare theorem assumes that a Walrasian equilibrium exists and that there are neither externalities from consumption and production nor firms able to affect the prices of the goods they sell. The Coase theorem assumes zero transaction costs. The Modigliani-Miller theorems assume that investors and firms can offer financial securities and/or transact securities on identical terms. The Vickrey-Myerson revenue equivalence theorem assumes that bidder values are independently and identically distributed and that they have no other information. These important theorems also assume that agents are all perfectly rational and, when they need to form expectations about future prices or other agents' strategies, their expectations are exactly correct!

Why do economists rely on such unrealistic assumptions? It is because a well-chosen simplification can remove the dust and smoke that obscures our view of the workings of economic forces. Although we celebrate the resulting theorems for the insights they deliver, we can apply them successfully only by being vigilant, working hard to understand not just the insights that simplified analyses provide but also how the designs and rule choices they inspire must be adapted to withstand the dust and smoke and much larger disturbances of the particular worlds in which the mechanisms will operate.

When I began my economics research career under the tutelage of my co-laureate, Robert Wilson, game theory had only begun to penetrate economics. Earlier economic theory focused on prices that cleared markets but had little to say about how such prices could be discovered. Wilson suggested that because auctions were a significant mechanism for simultaneously determining an allocation and discovering prices, and since their rules were often explicit, studying auctions would be a good way to begin understanding how prices and allocations emerge.

Much of the early auction research focused on questions related to the information content and the level of prices. How much of the bidders' information comes to be reflected in the prices that are paid? If a lot of information is reflected in prices, how does that affect the ability of bidders to profit? Do some auction rules lead to systematically higher expected prices than others?

The information problem took on a special urgency following the publication of an influential paper by Grossman and Stiglitz (1980). Their paper used an older, non-game-theoretic approach to model Fama's (1970) *efficient market hypothesis*, according to which asset prices reflect all information available to the market. In one version of their model, prices completely reflect *all* of the trader's information, including private information. Consequently, informed traders never profit from their information and so have no incentive to invest to acquire it. If that were right, however, then surely prices would not reflect such information, so a trader who acquired private information could profit by doing so. For auction theorists, this paradoxical finding just reinforced our belief that general equilibrium theory was the wrong platform to use for studying market clearing among investors with private information.

Wilson (1977) took a first stab at an alternative formulation in an auction model with "common values," meaning one in which there is an asset for sale that has the same value $\tilde{V}$ to all bidders. In the game theoretic auction model, there is no information paradox. Prices reflect bidders' information because they depend on the bids and bidders acquire information to the extent that they can profit from it.

Wilson's model famously highlighted the *winner's curse*, which arises because a winning bidder tends to be someone who has overestimated the value of the asset. To formalize that conclusion, suppose that each bidder $n = 1, \dots, N$ makes an independent, unbiased estimate X_n of the value $\tilde{V}$ (so $\mathbb{E}[X_n | \tilde{V}] = \tilde{V}$) and uses some increasing bid strategy $b_n(X_n)$. Bidder n wins only if its bid exceeds the highest opposing bid $\omega_n = \max_{j \neq n} b_j(X_j)$. Conditional on $\tilde{V}$, the estimates are independent, from which it follows that $\mathbb{E}\left[X_n \middle| \tilde{V}, X_n > b_n^{-1}(\omega_n)\right] \geq \mathbb{E}[X_n | \tilde{V}] = \tilde{V}$. According to this inequality, if Bidder n wins, one can infer that its estimate is larger, in expectation, than the asset value. A rational bidder needs to account for that in deciding how much to bid.

Despite the fame of the winner's curse, the main thrust of Wilson's paper was different. He sought to determine whether the equilibrium auction *price* might provide a good estimate of the common value $\tilde{V}$. For that analysis, suppose that N bidders participate in the auction and let $b_N^*(\cdot)$ be the Nash equilibrium bidding function that maps estimates into bids. The rules specify that the highest bid wins and the winner pays the amount of his winning bid. Let $P_N^* = b_N^*(\max(X_1, \ldots, X_N))$ be the winning bid. Wilson's paper identified a sufficient condition on the distributions to imply that, $\underset{N\to\infty}{\text{plim}}\, P_N^* = \tilde{V}$: the equilibrium price converges in probability to the actual value as the number of bidders grows.

My first published economics paper (Milgrom, 1979) extended Wilson's analysis to identify a necessary and sufficient condition for his conclusion. My condition was a purely statistical one, namely, that for all $v \in supp(\tilde{V})$ $v \inf_S \frac{\Pr\{X_n \in S | \tilde{V} < v\}}{\Pr\{X_n \in S | \tilde{V} = v\}} = 0$. Informally, this condition says that when the value is actually v, there is some information a bidder can acquire to make him nearly certain that the value is at least v. I used a statistical argument to show that for *any* functions b_N, the conclusion $\underset{N\to\infty}{\text{plim}} \max_{n \leq N} b_N(X_n) = \tilde{V}$ can hold *only if* my condition is satisfied. For the converse, I proved properties of the game-theoretic equilibrium strategies to show that my condition implies $\underset{N\to\infty}{\text{plim}}\, P_N^* = \tilde{V}$.

My second published paper continued the investigation of the information content of prices but shifted the focus from first-price auctions to second-price auctions, allowing bidders to decide what information to gather. In a second-price auction, since the price depends only on the highest losing bid, a winning bidder who has a better value estimate than the second highest among the others can earn a positive profit from the auction. This analysis showed promise by avoiding the Grossman-Stiglitz paradox, but it was limited to one-sided auction markets in which the bidders are just the buyers.

To analyze securities trading, I needed to study markets in which both buyers and sellers were concerned about the winner's curse. Milgrom and Stokey (1982) took a first step in that direction. It modeled trade among risk-averse agents with strictly concave utility functions u_n, each with some private information X_n about the values of various possible trades and some random endowment Z_n. The agents engage voluntarily in some mechanism that leads to a trade delivering a value increment to agent n denoted by $\tilde{V}_n$. Feasibility of the trade implies that $\sum_n \tilde{V}_n = 0$.

Let T denote the set of information profiles according to which all agents agree to trade using the mechanism. To avoid the winner's curse, rational agents account for others' willingness to trade, so n should participate only if $\mathbb{E}\left[u_n\left(\tilde{V}_n + Z_n\right) - u_n(Z_n)\,\middle|X_n, X \in T\right] \geq 0$. All agents agree to use the mechanism only when $T \subseteq \left\{X\middle|(\forall n)\mathbb{E}\left[u_n\left(\tilde{V}_n + Z_n\right) - u_n(Z_n)\middle|X_n, X \in T\right] \geq 0\right\}$. Can differences in information alone lead to trade? To answer that question, we assumed it was common knowledge that the initial allocation Z is Pareto optimal, that is, that with no information, no trade could increase every agent's expected utility. According to our "No Trade Theorem," these assumptions imply that for all n, $\tilde{V}_n 1_T = 0$, where 1_T is the indicator of the event that trade occurs: "Risk averse traders never make (non-zero) trades based solely on differences in information."

Our analysis implied that for trade to take place among rational traders, there must be at least the possibility that some Pareto-improving trade might exist. That is, there must be a *transactional motive* for trading – not just a speculative motive. As typical examples of a transactional motive, an agent may want to buy securities because she has just added savings to her retirement account, or to sell securities in order to pay this year's college tuition.

Once some traders can have transactional reasons to trade, speculators and others with information about $\tilde{V}$ may find themselves with profitable speculative opportunities. Glosten and Milgrom (1985) studied that by assuming that any trader arriving at the market may have a transaction motive or a speculative motive or a mix of both. We studied how a trader's information comes to be reflected in prices and whether the price a trader pays already reflects some of his own information.

In our model, a group of "specialists" compete against one another to serve customers who come to market one by one to trade in some security. For the specialists, each unit of the security can be sold tomorrow for some value $\tilde{V}$. The specialists set *bid* and *ask* prices today: a bid is an offer to buy from a customer who seeks to sell; an ask is an offer to sell to a customer who seeks to buy. The customer sees all the bid and ask prices and takes the best deal. Price competition among the specialists drives their profits to zero on every buy and sell transaction. When a *Trader* arrives at time t, the specialists' zero-profit conditions are that $Bid_t = \mathbb{E}_t[\tilde{V}|Trader\ buys\ at\ Bid_t]$ and $Ask_t = \mathbb{E}_t[\tilde{V}|Trader\ sells\ at\ Ask_t]$, where $\mathbb{E}_t$ is the expectation conditional on the public information available just before the Trade acts at time t. If the Trader sells or buys at time t, the actual transaction price P_t is equal to either Bid_t or Ask_t, respectively. Let $\mathcal{F}_t$ denote the specialists' information just after the Trader makes its choice at time t. If any transaction occurs, then $P_t = \mathbb{E}[\tilde{V}|\mathcal{F}_t]$. This formula implies that the stochastic process

$(P_t, \mathcal{F}_t; t \geq 0)$ is a martingale: prices have no predictable tendency to drift upward or downward.

Most importantly, this model resolves the Grossman-Stiglitz paradox while still incorporating a version of Fama's efficient markets hypothesis. In the model, transaction prices at any time t are equal to $\mathbb{E}[\tilde{V}|\mathcal{F}_t]$ – the expected value of the security conditional on the available information $\mathcal{F}_t$. This is the so-called "weak form" efficient markets hypothesis. What's more, the price P_t at which the Trader transacts does incorporate *some* of its own pre-trade information, namely, whether it had decided to buy or to sell. However, unlike the Grossman-Stiglitz model, the price does not generally incorporate all of the trader's information: a trader with sufficiently accurate information about $\tilde{V}$ may be able to profit.

The Glosten-Milgrom paper became one of my most highly cited ones. Along with the paper by Kyle (1985), it helped to launch the literature on financial market microstructure, which examines the detailed rules of trading and how those affect not only the informational properties of prices but also market efficiency, the size of the spread between bid and ask prices, and other market features.

How Auction Rules Affect the Bidders' Payoffs and the Seller's Revenue

Another set of questions that attracted economists' attention in my first era of auction research concerned whether and how bidders' payoffs and the seller's expected revenue depended on the auction rules. When I began my work, there were two main game-theoretic models for studying auctions: the Vickrey model and the Wilson model. My frequent collaborator Robert Weber and I introduced a new model that subsumed those two formulations, developed new methods to study payoff questions, and provided answers to the payoff and revenue questions.

In Vickrey's "private values" model, there are $N \geq 2$ bidders with values v_n that are independently and identically distributed according to some continuous distribution F on an interval $[0, \overline{v}]$. Bidder n knows v_n but does not know the others' values. A bidding strategy $b(\cdot)$ is a mapping taking the bidder's value v_n into a non-negative bid $b(v_n)$. The price p that the winning bidder pays depends on the auction rules and may depend on all of the bids. If bidder n wins the auction, his payoff will be $v_n - p$; if he loses, he pays zero and has a zero payoff.

Wilson's model, motivated by the case of oil companies bidding for drilling rights on tracts leased by the federal or state governments, also assumed that there are $N \geq 2$ bidders. Since the value to any bidder was determined mainly by the amount of oil that could be extracted as well as its depth and the cost of extraction, Wilson assumed that the value $\tilde{V}$ was about the same to all bidders, but that different companies' geologists and engineers made different estimates X_n of that value. If a bidder acquires the item for a price p, it earns a payoff of $\tilde{V} - p$, while other bidders receive zero payoffs. Wilson assumed that, conditional on $\tilde{V}$, the different bidders' estimates were statistically independent. As we have seen above, if bids are increasing functions of one's own value estimate, this necessarily implies a winner's curse.

The game-theoretic analyses in Vickrey (1961) and Vickrey (1962) gave birth to modern auction theory. In those papers, Vickrey analyzed Dutch descending auctions, English ascending auctions, and sealed tenders, which economists came to call "first-price" auctions. Each was modeled as a game with its own set of rules.

In a Dutch descending auction, the price starts high and declines almost continuously until some bidder shouts "Mine" and wins the item at that price. As Vickrey observed, a ("reduced") strategy in that auction is described by a single number or "bid": the price at which the bidder plans to shout "Mine." The bidder making the highest bid is the first to shout "Mine" and becomes the winner at that price. As Vickrey observed, the same mapping from bids to outcomes describes the first-price auction, so the two auctions are *strategically equivalent*. Thus, using Nash equilibrium for predictions, they lead to the same predicted outcomes.

In an English ascending auction, prices start low and rise as long as two or more bidders are active. As modeled by Vickrey, no bidder knows how many others are active, so a strategy is again characterized by a single price at which the bidder will become inactive. Let's call that single price a "bid." The participant with the highest bid wins, but the price is equal to the second-highest bid. The same mapping describes the rules of what is now called the "second-price" auction, so again the games are strategically equivalent. In those games, it is a dominant strategy for a bidder whose value is v_n to set his bid equal to v_n.

Vickrey worked out the equilibrium strategies and computed the expected prices and found a surprise: even though the first- and second-price auctions involve different rules, prices and strategies, the expected equilibrium price is the same for both! More precisely, let $b(\cdot)$ be the symmetric equilibrium strategy of the first-price auction and let $v^{(j)}$ denote the j^{th} highest among the values $v_1, \dots, v_N$. Vickrey's computations established that $\mathbb{E}\left[b\left(v^{(1)}\right)\right] = \mathbb{E}\left[v^{(2)}\right]$. This finding remained a puzzle for two decades until Myerson (1981) provided a further generalization: any two auction mechanisms that always lead to the same equilibrium assignment of the good, while requiring no payment by losers, necessarily lead to the same expected equilibrium price: the *Revenue Equivalence Theorem*.

In Milgrom (2004), I developed a new approach to extend the logic of the Revenue Equivalence Theorem and to derive new results.

My approach relied on the envelope theorem of Milgrom and Segal (2002). As applied to Vickrey's private value model, it works like this.

When a bidder places a bid (or more generally makes some moves or plays some strategy) σ in any mechanism, his choice leads to some probability of winning $\alpha_n(\sigma)$ and incurs some expected payment $\beta_n(\sigma)$. Letting $\sigma_n(v_n)$ denote an optimal bid for the bidder as a function of his value v_n. The bidder's maximum expected profit, $\pi_n(v_n)$, is defined by:

$$\pi_n(v_n) \stackrel{\text{def}}{=} \max_{\sigma} v_n\alpha_n(\sigma) - \beta_n(\sigma) = v_n\alpha_n\big(\sigma_n(v_n)\big) - \beta_n\big(\sigma_n(v_n)\big).$$

According to the "integral form" envelope theorem, $\pi_n(v_n) = \pi_n(0) + \int_0^{v_n} \alpha_n\big(\sigma_n(s)\big)ds$. If, as in many auctions, a bidder with value zero never wins or makes any payment, then $\pi_n(0) = 0$, so $\pi_n(v_n) = \int_0^{v_n} \alpha_n\big(\sigma_n(s)\big)ds$. In this case, the expected profit of a bidder with any value or "type" v_n depends only on the probabilities of winning $\alpha_n\big(\sigma_n(s)\big)$ for the lower types, that is, for $s < v_n$. This formula implies my *Payoff Equivalence Theorem:* if the bidders' types are statistically independent and if in the equilibrium of two mechanisms, γ and $\hat{\gamma}$, the bidder with the highest type always wins, then $\alpha_n^{\gamma}\left(\sigma_n^{\gamma}(v_n)\right) = \alpha_n^{\hat{\gamma}}\left(\sigma_n^{\hat{\gamma}}(v_n)\right) = F^{N-1}(v_n)$, so for both mechanisms, $\pi_n(v_n) = \int_0^{v_n} F^{N-1}(s)ds$. One implication of this formula is that every type of every bidder has the same expected payoff in both mechanisms. Since the expected value of the allocation is also the same for both mechanisms, the expected revenues are equal, too.

Milgrom and Weber (1982) used a closely related approach to analyze an auction model that subsumed the Vickrey and Wilson models. This model extended the Vickrey model in two ways: by allowing each bidder's value to depend on other bidders' information and by replacing the assumption of statistically independent types with a weaker assumption, described below. The envelope theorem approach, implicit in the Milgrom-Weber analysis, allowed us to prove revenue equivalence in a larger set of models and, for our most general model, to prove weak inequalities relating payoffs and revenues in different auction designs.

Because a bidder's information in our model could be different from his value, we need additional notation to describe it. In this section, bidder *n*'s information will be called its *type* and denoted by $t_n \in \mathbb{R}$. The bidders expected value of the item conditional on the full profile of types is denoted by $V_n = V(t_n, t_{-n})$. We assumed that $V(\cdot)$, which is a function of N arguments, is increasing in its first argument and symmetric and non-decreasing in its remaining $N - 1$ arguments. Vickrey's formulation is the special case in which $V_n = t_n$; Wilson's is the special case with $t_n = X_n$ and $V(t) = \mathbb{E}[\tilde{V}|t]$. For our initial analysis, suppose that the types are independent and identically distributed. Then, the value function

defined by $v(t_n) \stackrel{\text{def}}{=} \mathbb{E}[V(t_n, t_{-n})|t_j \le t_n \text{ for all } j]$ is increasing in t_n.
In auction game γ, the bidder's maximum expected profit is:

$$\pi_{\gamma n}(t_n) = \max_{\sigma} \int \Big(V(t_n, t_{-n})\alpha_{\gamma n}(\sigma, t_{-n}) - \beta_{\gamma n}(\sigma, t_{-n})\Big)\, dF(t_{-n}|t_n).$$

Suppose that, for auction game γ, there is a Nash equilibrium in which every bidder uses the same strategy σ_γ, which is an increasing function of the bidder's type. Invoking symmetry, we drop the subscript n from the functions $\alpha_{\gamma n}$, $\beta_{\gamma n}$ and $\pi_{\gamma n}$. Define $G^\gamma(\sigma, t_n) \stackrel{\text{def}}{=} \int \beta_\gamma(\sigma, t_{-n})dF(t_{-n}|t_n)$. The function G^γ expresses the bidder's expected payment as a function of its bid and its type. When the types are statistically independent, $(t_{-n}|t_n)$ does not depend on t_n, so the partial derivative is $G_2^\gamma(\sigma,\tau) \stackrel{\text{def}}{=} \partial G^\gamma/\partial\tau = 0$.

If equilibrium involves each bidder playing the increasing strategy σ_γ, then the winner is the bidder with the highest type, which leads to this helpful identity:

$$\int \alpha_\gamma\big(\sigma_\gamma(t_n), t_{-n}\big)dF(t_{-n}|t_n) = \Pr\Big\{t_n > \max_{j\neq n} t_j \,\Big|\, t_n\Big\}.$$

Applying the envelope theorem to study n's maximum expected profit function, we have:

$$\pi_\gamma'(\tau) = \frac{\partial}{\partial\tau}\Big(\mathbb{E}[V(\tau, t_{-n})|t_n = \tau]\Pr\{t_n > \max_{j\neq n} t_j \,|t_n = \tau\}\Big) - G_2^\gamma(\sigma_\gamma(\tau), \tau).$$

As previously observed, if the types are statistically independent, then $G_2^\gamma = 0$, and then no γ appears on the right-hand-side, that is, π_γ' *is the same for all such games*. If we also have that $\pi_\gamma(0) = 0$, then all games in this class have the same profit function π. This subsumes and generalizes the Payoff Equivalence Theorem.

The second way in which Milgrom and Weber generalized the theory was by allowing the bidders' types to be correlated in a particular way. Let $f(t)$ denote the joint density of the types. This model introduced the assumption that types are "affiliated," which means that $\frac{\partial^2 \ln f(t)}{\partial t_n \partial t_m} \ge 0$ for all t and all $n \neq m$. As a leading example, this condition is satisfied when there is a real-valued parameter θ with any prior density g such that, conditional of θ, each type t_n is independently distributed with conditional density $f_n(t_n|\theta)$ that satisfies the monotone likelihood ratio property. To prove that this example works, one may use the fact that the monotone likelihood ratio property can be expressed as $\frac{\partial^2 \ln f_n(t_n|\theta)}{\partial t_n \partial\theta} \ge 0$. The Payoff Equivalence conclusion does not generally apply without statistically independent types, but it can be replaced in symmetric models by the Milgrom-Weber *Linkage Principle*, which compares the expected profits and payments for all bidder types in a symmetric

Nash equilibrium in certain pairs of auction games. The intuitive idea of the principle is that any bidder's expected payoff, or "information rents," are lower if G_2^{γ} is higher, that is, if conditional on its expected payment, payments increase faster as a function of the bidder's actual type. For example, let us compare a first-price auction γ, in which $G_2^{\gamma} = 0$ (because the winning bidder pays its bid), to a second-price auction $\hat{\gamma}$, in which the winning bidder pays the second highest bid, the expectation of which can depend on the winning bidder's type. If types are affiliated, then given whatever bid the winner may make, $G_2^{\hat{\gamma}} \geq 0$: the conditional expectation of the second highest bid is an increasing function of the winner's type. This statistical linkage leads the winner's expected price to rise faster as a function in $\hat{\gamma}$, and one can use this and the envelope formula to conclude that $\pi_{\gamma} \geq \pi_{\hat{\gamma}}$: bidder profits are lower and seller revenues are higher in the second-price auction.

The Linkage Principle allows us to derive inequalities by focusing on the partial derivative G_2^{γ} in the envelope formula. Recall that this derivative is always zero when types are statistically independent, so inequalities among auctions arise from the statistical *linkages* between a bidder's type and the price that bidder pays. In the Milgrom-Weber "clock-auction" model of the English auction, the bidders all watch as posted prices rise continuously on a clock. Bidders are assumed to observe the price levels at which each bidder stops bidding and, in equilibrium, they bid up to higher prices when others are seen to bid more, so final prices are an increasing function of all the losing bidders' types. Given affiliation, all of those types tend to be higher when the winning bidder's type is higher, so G_2^{γ} is larger higher for the English auction than for the second-price auction. In another example, suppose that the seller has information t_0 about the value and that the entire random vector $t = (t_0, t_1 \dots, t_N)$ is affiliated. If the seller announces t_0, then the price depends on that announcement, which is positively related to each t_n. The technical work in this paper involves establishing that these positive linkages are of the right sort to increase G_2^{γ} as needed. Applying the Linkage Principle to the unique increasing symmetric Nash equilibrium in various games, we found these results:

- The expected Nash equilibrium price is higher in a second-price auction than in a first-price auction.
- The expected Nash equilibrium price is higher in an English ascending auction than in a second-price auction.
- The expected Nash equilibrium price is higher in each of the first-price, second-price, and English auction when the seller adopts a policy of revealing its information t_0 to all the bidders than when its policy is to withhold that information.

Allocating Radio Spectrum Licenses and the SMR Auction Design

The second era of my auction research launched in 1993, when the US government decided to change the way that it allocated radio spectrum rights. Until the mid-1980s, the rate of arrival of new applications to use radio frequencies was modest. When a dispute arose about the competing uses of some band of radio frequencies, the Federal Communications Commission (FCC) conducted comparative hearings to determine which use best served the public interest and awarded spectrum rights accordingly. It was an awkward and time-consuming process, which some favored because it enabled the regulator to press applicants to serve various interests and enhanced the power of the regulators and politicians. This system of comparative hearings was overwhelmed by the emergence of mobile phone technology in the early 1980s, which led to hundreds of applications to provide services to some small areas. Congress authorized the FCC to abandon comparative hearings and replace them with a lottery system, but that system was problematic, too. Applicants needed no particular qualifications to participate in these lotteries and sometimes won valuable rights to supply cellular services. In the best case, they would quickly sell those rights to a real telephone company that could provide an actual phone service. In the worst case, they would hold those rights for years, depriving consumers of any service.

In 1993, Congress acted again, calling upon the FCC to eliminate the lotteries and instead conduct auctions to sell the license rights. According to the law, the primary goal of the auctions would be to promote an efficient and intensive use of radio spectrum. Among the secondary goals was to capture a portion of the value for the US Treasury.

The FCC believed that the best way to promote efficient and intensive service was to promote competition among telephone companies, including local carriers like Cincinnati Bell, regional companies like Bell Atlantic, national companies like AT&T, MCI and Sprint which then provided long-distance phone service, and new businesses. These would be complicated sales of a set of valuable assets. Having no experience running even small, simple auctions, the FCC handed the task of planning the auction program to a team of its economists led by Evan Kwerel. The team sought guidance in the academic auction literature and cited some of my papers in a published *Notice of Proposed Rulemaking*. Despite their citations, there was little in the existing game-theoretic auction literature that could usefully guide the detailed choices needed for this auction.

Long before the application of game theory to study auctions, Ronald Coase (1959) had advocated using auctions to assign spectrum rights. In what was to be his first advocacy of the now-famous Coase theorem, he argued that so long as rights are tradeable, private exchange will lead them to migrate into the right hands. Moreover, he argued, allocating rights according to willingness to pay in an auction was likely to lead to a good initial allocation, reducing the amount of trade that would be needed.

Coase and his followers have the idea that, somehow, markets find a way to migrate resources to their highest value use. Writing before the emergence of information economics and modern market design, he did not worry about inefficient bargaining or about how parties could identify the most efficient arrangements. Today, most economists would agree that assigning property rights to permit exchange can be an important step to facilitating efficient exchange, but often more is needed. To assess the effectiveness of an unorganized market and the need for careful market design, we ask additional questions. Is the resource allocation problem a simple one in which a nearly efficient outcome can be reached by a series of profitable bilateral exchanges? Or are complex, multilateral arrangements needed? If market design is needed, what market rules can facilitate and promote the necessary trades?

From the perspective of 1993, the biggest market design challenge was one of "price discovery." Before the auction, bidders could not know on what terms different licenses might be sold. According to the price discovery perspective, the job of the auction was to make that clear. Many of the so-called experts contributing to the debate took other perspectives. One leading proposal was to employ rules like those used by Sotheby's or other traditional auction houses, according to which the spectrum licenses are auction lots to be sold one-at-a-time in a pre-determined sequence. In practice, however, such a sequential design would create intractable problems for the bidders. For example, how much should a company hoping to set up a nationwide business bid for the first licenses offered in, say, Chicago and Los Angeles, when there is still no information about the cost of acquiring licenses to serve New York and Boston? What will happen to the bidder if it succeeds in buying licenses in some but not all of the cities needed for its business plan? If bidders guess wrong about future prices, bidding mistakes could lead to an inefficient use of the frequencies.

Another proposal would have addressed that problem by allowing nationwide service companies to make *combinatorial* bids, which offer a single price for a desired package of licenses, such as one covering the whole country. But how would such bids be compared with individual bids for smaller licenses? If the total initial bids of the small bidders was too low to win against the national bidder, would there be a procedure to allow them to raise all their bids to win? How could the small bidders be encouraged to coordinate so that their total bid was high enough to be winning?

Questions like those led both the FCC and several telephone companies to seek the advice of academic economists who had studied auctions. Robert Wilson and I were approached by Pacific Bell, a regional telephone company, while the FCC hired Professor John McMillan. As I pondered the challenges of the spectrum allocation problem, I was inspired by my experience bidding in "silent" auctions at charity events. In a typical such event, people donate things to be sold in an auction. For example, one person might donate cooking lessons; a second person might donate an evening with a celebrity; a third, a weekend at a privately owned ski chalet; another, a bottle of wine, and so on. The items or descriptions of them are put on tables in a large room and everything is for sale at once in ascending auctions. There is a pad of paper and a pencil in front of each item. Bidders would write their name or ID number and a price on the paper, subject to the restriction that the bid must exceed the price on the preceding line by some minimum increment. There is a deadline for bidding, commonly set just before food is served or another fund-raising activity begins. The simultaneous ascending design allows a bidder who is unsure which item to bid on to begin with her most preferred items and then switch to others if the preferred items become too expensive. This process eases the bidder's task of deciding which item to bid for and how much to bid.

In contrast, consider a problem reported in earlier simultaneous sealed-bid designs for mineral rights in the US and for radio spectrum in Australia. In those auctions, many items draw only very low bids, as bidders, seeking to keep the total prices they pay within some allotted budget, bid high on just a few items and low or zero on the others. Often, the bidders guess wrong about which ones they can win. That is much less problematic in silent charity auctions as participants watch the bid sheets, learn about prices, and identify the best potential deals before the auction is over. Another advantage of simultaneous ascending auctions is that as bidders sort themselves out, a single bidder could become the highest losing bidder for many items, raising the prices of each of those.

In my experience, this design had worked pretty well for the charities, as many bidders regarded their prices as a contribution to a good cause. However, I could also see that the silent auction rules were being gamed by clever bidders, and in a way that might become much worse in a high-stakes spectrum auction with billions of dollars at stake. Bidders in the silent charity auction engaged in a form of "sniping," which means waiting until the last moment to bid. A bidder who wanted to buy some item could sometimes keep the price low by refraining from bidding until the last few seconds before the bidding deadline. By bidding only at the end, it might deny other bidders who would be willing to compete any opportunity to react.

To retain the advantages of the silent auction design while eliminating the sniping problem, Robert Wilson and I proposed our simultaneous multiple round (SMR) design. It differed in three ways from the silent auction.

1. Instead of being run in continuous time with all bidders present on site, we structured the auction to be run remotely in a *series of discrete rounds*. In each round, there would be a minimum price for each item, which was set by adding a minimum increment to the best previous bid. No information would be shared during a round. After a round closes, information about bids would be reported to the bidders and time would be allowed for bidders to digest the information and plan new bids. Then, the next round would begin.
2. To prevent sniping on individual licenses, we added a *termination rule*, which says that the auction does not end until there is a round with no new bids for any license. In this way, when some bidder raises another's price for a license, the second bidder always has a chance to raise back or to switch and bid on another license. There are other important details of a design like this one, such as what information to report to bidders after each round and also the choice of licenses to be offered for sale.
3. To ensure adequate progress of the auctions, I invented what has become a standard rule category in auctions – the *activity rule*. The activity rule says a bidder cannot bid for a larger amount of spectrum rights in any round than it had bid in the previous round. This rule prevents a bidder from waiting to see what others were doing before making its own commitments, helps the auction to develop meaningful prices before bidders must make their final bids, and shortens what could otherwise be an untenably lengthy process.

Preston McAfee had offered a similar simultaneous ascending design. The main difference was that instead of our termination and activity rules, McAfee's design specified that bidding on individual licenses would close after some number of rounds with no new bids on that license. This rule was designed to deal with the same problems, but it could force bidders to decide prematurely whether to continue bidding for the most preferred license or to switch away to a less desirable one, which would soon become unavailable under the termination rule.

To convince the FCC that our proposed SMR auction design was easy for the bidders to use, completely specified, and implementable in practice, I arranged to have it coded using Excel spreadsheets. Each bidder had a spreadsheet on which to enter its bids, with code to check that the bids satisfied the activity rules and any other eligibility rules. The auctioneer had a spreadsheet, too, which in each round imported bids from the bidder spreadsheets, rechecked that the bids satisfied all the rules, processed the bids and exported the round results back to individual bidders' spreadsheets. Although the real implementation would require internet communication features and much greater security, the Excel sheets were a sufficient proof of concept to convince the staff to recommend the SMR design to FCC chairman Reed Hundt, who approved it and ordered it to be implemented.

The initial design ran smoothly without any major glitches and the FCC trumpeted its success. The design was welcomed by bidders, celebrated in the popular press, and copied by several other countries. One *New York Times* article even declared it to be "the greatest auction in history." Since the initial auction, more than $100 billion of spectrum sales in the US and more than $300 billion around the world have used some version of these rules. All this, despite the initial absence of any deep theoretical treatment of why such an auction should work.

That importance of missing theory was driven home to me shortly after FCC auction #4, when Australian regulators who were free market enthusiasts came to Washington. They proposed the idea of "postage stamp" licenses, meaning thin slices of spectrum covering tiny geographic areas, which phone companies could assemble any way they liked, allowing maximum flexibility for their business plans.

The Australian proposal idea highlighted one of the gravest dangers of the SMR auction design. A small bidder in this auction who takes minimum bid prices as given in each round might find itself becoming the highest bidder on a set of licenses before learning that the prices of other licenses will rise so high that its business plan for those is unprofitable. The bidder must then decide whether to buy those licenses anyway, losing money, or to acquire the smaller set of licenses it is currently winning, which may be just a subset of what a profitable business plan requires. This risk of encountering this conundrum has come to be called the *exposure problem*. Evidently, this risk is absent only when increasing the price of some licenses cannot lead a bidder to prefer to demand fewer other licenses. A demand function with that property is said to exhibit "gross substitutes."

Kelso and Crawford (1982) deserve much of the credit for the next analysis. They had studied the role of the gross substitutes condition in a model of labor markets that is quite a close cousin of my auction model, and the conclusions of the models are close cousins, too. In Milgrom (2000), I found that if bidders in the simultaneous multiple round auction bid straightforwardly as price takers – as small bidders are encouraged to do by the activity and termination rules – and if the licenses are substitutes for all bidders, then the final auction allocation is *approximately* efficient and its prices *approximately* clear the markets for all licenses. Here "approximately" means that the efficiency and market-clearing conclusions hold exactly for a different economy in which bidder values for some licenses are only slightly lower, with no value difference exceeding one bid increment.

Market Design for Substitutes: Auctions and Matching

Modern market design was born in 1994 with the introduction of the SMR auction design for radio spectrum and the redesign of the National Resident Matching Program for doctors based on the matching theory of Gale and Shapley (1962). The substitutes condition plays a central role in matching theory, too. During the Gale-Shapley algorithm, when a hospital rejects the offer of some doctor and later receives an offer from other doctors, the substitutes condition means that if the first doctor's offer were still available, the hospital would still reject it. At the end of the algorithm, any doctor who prefers to work for some hospital has proposed to it and been rejected and would still be rejected if she renewed her proposal. That is why the final outcome is a *stable matching*.

The Gale-Shapley matching algorithm seemed auction-like to me, at least from the standpoint of the hospitals. The algorithm is one in which the hospitals reject inferior offers and hold onto the currently best ones but remain open to rejecting those in favor of still better offers that may come later. That raises the question: How important was this similarity and how broadly could it be used for applications?

I collaborated with a graduate student to answer that in Hatfield and Milgrom (2005). In our model, a contract was a bilateral agreement between two parties that had three blanks to fill in: the name of the offering party or "buyer", the name of the counterparty or "seller", and any extra *terms*. Mathematically, it was a triple $(b, s, t) \in B \times S \times T$, where the three sets B, S, T are all finite. When T is a singleton, that is a standard matching model. To represent offers that are prices, one can set $T \subset \mathbb{R}_+$ and prescribe that the buyer always prefers lower prices and the seller higher prices. The terms could also have many other interpretations. For example, in the marriage problem, a proposal that says, "we'll move to Arkansas and take over dad's store when he retires next year" might elicit a different response than "we'll move to the beach in LA and go surfing every weekend." We showed that many of the results of matching theory continue to apply as long as offers are substitutes, and our framework also enabled new applications in which the substitutes condition does

The first step in the usual matching mechanism is for each side to describe its preferences. In a simple matching with contracts model, parties on one side – let's call them the "doctors" – seek out just one contract with one hospital, but those on the other side, let's call them the "hospitals" – are seeking to hire a collection of doctors. For the doctors, if all hospital-term pairs are acceptable, that means reporting a list of $|H \times T|$ hospital-term pairs. For a hospital that is seeking to report its preferences over subsets of contracts with many doctors, there are $(1 + |T|)^{|D|}$ combinations to consider, because each doctor can either be excluded from the hired set or have any of $|T|$ contracts. Even if we limit attention to sets with exactly k doctors, there are still $|T|^{\binom{|D|}{k}}$contract combinations to be ranked, which is often far too many for a simple list.

In Hatfield and Milgrom (2005), we begin the task of finding simple languages to express preferences for packages of items, especially when the terms are monetary and/or the doctors are substitutes for the hospitals. The first simplification comes from assuming quasi-linearity of preferences, so that preferences are described simply by specifying a value for each package. Even so, in a spectrum auction with $|D| = 1000$ items for sale, the number of subsets is 2^{1000}, and no list of such a length can be reported in practice. We suggest further restricting the set of values by imagining that each worker is to be assigned to a specific role and has a value in that role, with at most one worker in each role. Then, for a set of k workers assigned to r roles, it takes just kr numbers to describe the values of any set of workers. The same logic can be applied when we allow the number of workers to vary and when we suppose that the hospital already has some doctors under contract. We require just r values per worker to describe the *endowed assignment valuations*. All such valuations make the doctors into substitutes, so the usual theorems apply.

In Milgrom (2009), I introduced a closely related language to report substitutes preferences in auction problems. There were two main differences in this language. First, many auction problems feature large homogeneous categories of goods. For electricity, power from several sources might contribute to serve some location-time pair and there might be transmission constraints that limit acquiring certain amounts of power. Second, for simplicity, I eliminated the concept of roles and instead focused on bids to serve some purpose. Constraints reflected simplified transmission constraints, limiting the amount of power to be purchased from various power sources. I showed that so long as these constraints had a hierarchical ("laminar") form, the resulting demand satisfies the gross substitutes condition. I also showed that if the constraints were integer variables, then the optimal assignment would be integer as well, so the same setup could be used, for example, to allocate shipping containers, in which quantities are divisible but it is rarely economical to ship any quantities that do not correspond to a whole number of containers.

The Exposure Problem and Combinatorial Auctions

How big a problem is it in practice that bidders do not regard all items as substitutes? Are there strategies a bidder can use in an SMR to avoid the exposure problem? Good auction design takes close account of the perspectives and capabilities of the bidders.

In Bulow, Levin and Milgrom (2017), we recount the strategies that we used to advise a bidder – Comcast – in the auction for Advanced Wireless Services (AWS) conducted in 2006. That sale used an SMR auction design. The bidder instructed its auction team as follows: buy at least 20MHz of bandwidth coverage in every major US city, about 2/9 of all the spectrum for sale – but only if the total cost was less than the maximum budget that the bidder had allotted. Otherwise, buy no coverage at all.

According to received theory, those instructions create a classic exposure problem, in which our client might make winning bids on some licenses without being able to afford the whole set that its business plan required.

In practice, we concluded that the key to eliminating most of the exposure risk was to be able to forecast final total prices from the early bids. While prices were low, our client would bid for what it wanted. Later, if the forecasted total price was too high, the bidder would exit and plan to resell any licenses that it had acquired when prices were low. Our forecast, which proved to be accurate, indicated that the desired package of licenses would be affordable. The bidder relied on that assessment and successfully acquired its target package. But how was the forecast made?

In most FCC spectrum auctions, with billions of dollars at stake, budget constraints often limit the total amounts that the major bidders can buy. When those constraints are all binding, the total price paid for all licenses is approximately equal to the sum of the bidders' budgets. To estimate that sum from the early bidding behavior, we first compute for each round the money amounts that bidders risk spending if their bids become winning, which is the sum of the standing high bids from the previous round plus the new bids on other licenses at the current round. This *budget exposure* statistic rises steadily in the early rounds of an auction, as prices rise. At some point, however, budget exposure often seems to hit a ceiling as all bidders become budget-constrained and stop bidding for lower priority licenses. This ceiling was our estimate of the total budget amount committed to the auction by all bidders.

This bid advisory experience reminds us as designers how important it is to consider the perspective of the bidder and, in particular, cautions us not to exaggerate worst-case exposure problems in developing and assessing auction designs.

While the exposure problem is sometimes manageable for bidders, there can be other times when the exposure problem is especially problematic. In such cases, taming the problem requires adopting a *combinatorial* auction design, which refers to any auction design in which bids are offers to buy entire packages of items, rather than a collection of individual offers for the items in the package. Combinatorial designs face three new kinds of special challenges:

1. Bidders typically cannot bid on all the packages. For example, with just 100 items for sale, there are 2^{100} packages available, which is far too many to be included in a bid list. What are some useful compact ways to express bids on many packages without simply listing packages and values? Lacking such a compact expression, how can we design an auction to exchange information among bidders to guide them to bid on the correct packages?
2. If bidders do bid on a large number of packages, then the problem of finding the set of bids that maximize the total price is NP-hard. How should the auction system select winners if it is unable to optimize? What limits can be put on the package bids to ensure that the system can solve the problem?
3. How should prices be set? Economic theorists and computer scientists are often quick to point to the Vickrey auction rule, but it requires optimization, which can make its prices intractable to compute. Also, as Ausubel and Milgrom (2006) showed, Vickrey prices can be uncompetitively low and encourage collusive strategies, among a longer list of problems. Are there alternative pricing rules that mitigate these problems?

For the first challenge, Ausubel, Cramton and Milgrom (2006) introduced the combinatorial clock auction (CCA) as a design to communicate relevant price and package information. In that design, bidders would decide which items to include in their packages and the system would determine prices using rules similar to the SMR auction, in which prices are increased for any individual items included in the demanded package by more than one bidder. Once total demands at the posted prices were no more than the supply, this "clock stage" of the auction would end. Bidders would have one more chance to submit a limited number of additional package bids that are consistent with the package values expressed during the clock stage. The precise notion of "consistent" has gone through several versions over time. The main advantage of this design was to guide bidders to packages that fit well with those expressed by other bidders.

Given that bidders are unable to bid for all packages, which are the important packages on which bids must be received for the auction to perform well? Which packages are real bidders likely to identify for bidding? Kagel, Lien and Milgrom (2010, 2014) studied those questions. We characterized sets of "efficiency relevant packages" and "core relevant packages" and found that, to generate efficient outcomes or competitive prices, it was sufficient that bidders bid aggressively on those packages, even excluding any bids on other packages. In addition, we found that for some clock-guided designs, bidders in an economics laboratory were guided by prices to bid in each round mainly for the packages that appeared to be most profitable at the current clock prices.

My study of bidding languages to express bids compactly began in Hatfield and Milgrom (2005) and Milgrom (2009). Those two papers offered intuitive ways to describe a large class of substitutes values, but these languages do not describe any complementarities. Eilat and Milgrom (2011) offers a language to describe simple economies of scale and scope that arise from fixed costs while Bichler, Milgrom and Schwarz (2020) adds a capability to express general economies of scale in a product category and limited economies of scope across categories.

The third combinatorial issue is how to set prices from a set of combinatorial bids. Day and Milgrom (2008, 2013) studied *core-selecting auctions*, which are auctions that satisfy three conditions:

1. the winning bids comprise the feasible set that maximizes the total bid,
2. no bidder pays a price higher than its own bid, and
3. the total auction revenue is not less than another set of bidders has offered to overturn the allocation.

The Vickrey auction, which eliminates bidders' incentives to misreport, is not a core-selecting auction, so we characterized the auctions among those that are core-selecting that minimize the total incentive to misreport, which we called the *minimum-revenue core-selecting auctions*. A minimum-revenue core-selecting pricing rule was combined with our CCA communication design to create another one of the most widely used auction rules for radio spectrum sales.

Day and Milgrom (2013) also included a conceptual breakthrough, highlighted in point 3 above. In contrast to earlier papers in mechanism design theory, which had treated incentives as a constraint that must always be exactly satisfied, our analysis relaxed that constraint in favor of a constraint on total auction revenue. We instead treated the incentive to deviate as an objective to be minimized.

Mixing Theory and Computations

Motivated in part by my earlier finding that theoretical concerns like the exposure problem are sometimes exaggerated, my recent work incorporates computational studies, as well as formal theorems, into the analysis of auction mechanisms. Two examples of this approach are discussed below. The first arose from an auction design problem for internet display advertising. Events have since overtaken that design – it no longer addresses the most important challenges of that industry – but my coauthors in that work deserve credit for breaking ground in combining theory and computations. The second example is perhaps the capstone of my career: the Broadcast Incentive Auction. Auctionomics' design recommendations for that auction were adopted by the Federal Communications Commission for what appears to be the most complex resource reallocation problem in history. Working with Ilya Segal and Kevin Leyton-Brown, we devised new theory and software algorithms to create a new auction design that was simple and intuitive for bidders and could run that multi-faceted process smoothly and successfully.

Adverse Selection in Internet Advertising

My co-laureate, Robert Wilson, had introduced an auction model that emphasized the winner's curse, which is a form of adverse selection. After Google's sponsored-search auctions proved that targeted, intention-based advertising could be hugely profitable, internet sites that had been using traditional contracts to sell to traditional advertisers, who seek to enhance their brands or inform consumers about special events, began to introduce auction methods to reach a new category of advertisers. These new advertisers look at consumers individually, measuring and predicting some aspect of ad performance such as whether consumers click links, fill forms, or buy products. Auctioning targeted ads, however, creates a problem of adverse selection for brand advertisers, because each ad is shown to just the remaining viewers that no performance advertiser thought was worth paying for, instead of reaching a representative mix. The remaining impressions are, on average, much less valuable.

At the time, second-price auctions were most commonly used to sell such advertising, due to their efficiency and dominant strategy properties. I wondered: Is there a way to use an auction to sell impressions that protects the contract advertiser from adverse selection but still allows an advertiser who is an especially good match for some customer to select that customer for targeting? What would such an auction look like?

Writing with two of my students in Arnosti, Beck and Milgrom (2016), we found a surprising affirmative answer to the first question using a model that distinguishes brand advertisers from performance advertisers. For the brand advertiser whose ad would, in the absence of performance advertisers, be the one that is shown to the visitor, the value of that ad impression is X_0V, which is assumed to be unknown to the brand advertiser (and is the reason the brand advertiser does not engage in targeting). For any performance advertiser $n = 1, \dots, N$, the value is X_nV, which that advertiser knows. The common V term is a measure of valuable general characteristics, such as the visitor's income and how engaged she is with that site and how responsive she may be to online advertising. The X_n terms measure the match quality. For example, a consumer who has recently been reading about and shopping for home loans will be particularly valuable to a retail mortgage company like Quicken Loans. The key assumption of the model is that the two random variables X_0 and V and the random vector $X = (X_1 \dots, X_N)$ are statistically independent. This means that the vector of performance advertiser match qualities X is uninformative about both the brand advertiser's value X_0V and the common term V. In a direct mechanism, the performance advertisers make their reports XV. Let B denote the set of bid profiles that result in an award to the brand advertiser. If $XV \notin B$, then the impression is assigned to some performance advertiser.

Our paper originated a new approach to auction market design that mixes axiomatic and computational approaches in a promising recipe. The axioms were these: One requires that the mechanism is *strategy-proof* for any fixed number of bidders, so each bidder's dominant strategy is to report truthfully: $b_n = VX_n$. *Deterministic mechanisms* and *anonymity* axioms then imply that only the performance advertiser that places the highest bid can ever become the winner. Another axiom requires that the auction mechanism be *false-name proof*, meaning that if there are at least two bidders, no bidder can influence the price or allocation by submitting an additional low bid under a false name. Next, the allocation rule must be *adverse-selection free*, which means that the probability that the impression is awarded to the brand advertiser is independent of V and X_0V which means that $X \in B \Leftrightarrow (\forall v) vX \in B$.

For notational simplicity, let us label the performance bidders in the order of their submitted bids, from largest to smallest, so $b_1 > b_2 > \cdots$. The false-name proof axiom implies that an allocation to the brand advertiser depends only on the two highest bids (b_1, b_2), and adverse-selection free then implies that it depends only on the ratio of values

$\frac{X_1}{X_2} = \frac{VX_1}{VX_2} = \frac{b_1}{b_2}$. The allocation rule of a strategy-proof mechanism is monotonic, so using the preceding there exists some $\beta \geq 1$ such that bidder 1 wins if and only if $b_1 > \beta b_2$. To be strategy-proof, the auction mechanism must also set the winning bidder's price equal to its minimum winning threshold βb_2. Our main theorem in the axiomatic section summarizes all this as follows: the set of auction mechanisms that satisfy the listed axioms forms a parameter class indexed by $\beta \geq 1$, as just described. By varying β, the probability that any impression is assigned to the brand advertiser can be anything between zero and one. Setting $\beta = 1$ recovers the Vickrey auction among performance bidders only, excluding the brand advertiser.

For the computational part of our analysis, we need additional assumptions. We assume that $X_1, \ldots X_N$ are independent and identically distributed according to a power law distribution: for any $x \geq 1$, $\Pr\{X_n > x\} = x^{-\rho}$, where $\rho > 0$ is a parameter of the distribution. We set X_0 equal to a constant, allow V to be drawn from any distribution G with finite expectation, and vary β to find the mechanism in our class that maximizes efficiency. To measure efficiency, we divide the expected value under our mechanism by the maximum expected value of any allocation rule and find that the worst-case performance of this mechanism is surprisingly good: $\min_{\rho, X_0, G, N} \max_{\beta} \text{EfficiencyRatio} = .948$.

Intuitively, the high efficiency ratio means that the new mechanism assigns the impression to performance advertisers whenever the gains from doing so are substantial. This is a characteristic of fat-tailed distributions like the power law, for which a large fraction of the total gain from assigning impressions efficiently to performance advertisers is realized when X_1 is high. For fat-tailed distributions, when X_1 is high, $\frac{X_1}{X_2}$ tends to be high as well, so the new auctions tend to capture much of the gains.

The Broadcast Incentive Auction

After Apple's introduction of the iPhone in 2007, consumer demand for wireless data exploded. Soon after that, the FCC began to explore ways to phase out older, less valuable uses of the spectrum to make use for more valuable ones. One of the first targets was the spectrum allocated to over-the-air television broadcast, which had declined in relative value as consumers increasingly used alternatives like cable, satellite, and the internet to watch their favorite television shows.

The TV spectrum reallocation problem was made complicated beyond any precedent by several factors.

First, the effective engineering of the system requires that the same channels be assigned to television broadcasting across the United States, and the engineering works best if the channels are also the same in Canada and Mexico, so that there is no problem of broadcast interference near international borders. Clearing channel 42 in Chicago is hardly worthwhile unless one can clear the same channel in New York and Los Angeles. Moreover, if channel 42 is used for broadcast television, then channels 41 and 43 cannot be used effectively for mobile data. To create maximum value, the reallocation and reassignment of stations to new channels to free a good set of channels for mobile data needed to be well coordinated across the whole country and even internationally.

Second, the amount of spectrum to be cleared was not known in advance. Some TV station owners would be willing to sell their broadcast rights for a suitable price and mobile telephone companies would pay something to buy, but how much could be traded and at what cost? FCC staffers Evan Kwerel and John Williams suggested that a market process would be the best way to answer those questions.

Third, there were about 2400 TV stations, each with different conditions of supply and demand that may require setting its own distinct market price. On the supply side, TV station values depend mostly on the viewers they reach, and there are huge variations among stations in the numbers and incomes of their viewers. For example, there are many times the number of viewers of a full power station in New York than of stations in Tampa or Boise or Stockton, and average incomes vary in those cities, too. On the demand side, because every station location has a different coverage area, it also has a different role in avoiding interference with other TV operations, giving it a distinct value to the FCC.

Fourth, determining whether a set of channels can be cleared using a particular set of stations is extremely challenging. There is a branch of computer science called “complexity theory,” which classifies the difficulty of computational problems. To evaluate the difficulty of the present problem, it is helpful to think of each TV station as a node in a graph. Let us say that two TV stations are connected by an edge in the graph if those stations are too close to share the same channel. More precisely, they are connected by an edge if assigning the same channel to both stations would cause more than 0.5% of either of their viewing audiences to suffer signal impairment from the resulting interference. With this framing, the question of whether it is possible to assign a channel from a given set of channels so that no two connected stations share a channel is an instance of the famous graph-coloring problem: is it possible to assign a color from a certain set of colors to each node so that no two connected nodes are the same color?

This graph coloring problem is known to be NP-complete, which means that for any known algorithm, the worst-case solution time grows exponentially in the problem size. The North American interference graph, with about 130,000 edges, is a large problem, and numbers like $1.1^{130{,}000}$ are almost unimaginably large. That difficulty is just to check whether purchasing a certain set of stations makes clearing a set of channels feasible. Finding the solution that minimizes total cost is, in practice, even harder. Any practical market algorithm would be unable to make use of optimization, so traditional combinatorial auction designs could not be used.

The auction design team's solutions to the full set of auction design challenges are spelled out in Milgrom and Segal (2020) and Leyton-Brown and Milgrom and Segal (2017), but I will focus here just on the problem of designing a "reverse" auction to purchase TV broadcast rights. The first step in creating that design was to introduce a vast new collection of descending clock auction designs, which are formalized as follows.

There are N bidders and a finite set $\mathcal{P}$ of possible prices. In each round t of the auction, each bidder n is presented with a price $p_n^t \in \mathcal{P}$. If the price is different from that of the preceding round ($t \geq 2$ and $p_n^t \neq p_n^{t-1}$), then the bidder may reject the price and exit from the auction, making no further accept-reject decisions. Let h_t be a list of the reject decisions made at round $t = 1,2 \ldots$. The history of such reject decisions through round t is denoted by $h^t = (h^{t-1}, h_t)$ where $h^0 = \emptyset$. Let $\mathcal{H}$ be the set of possible histories.

A descending clock auction is described by any function $p: \mathcal{H} \to \mathcal{P}^N$ satisfying $p(h^t) \leq p(h^{t-1})$ for every history $h^t \in \mathcal{H}$. The interpretation is that $p(h^{t-1})$ is the vector of prices offered to the bidders in round t and that the auction ends after round T if $p(h^T) = p(h^{T-1})$. At that time, winners and losers are determined. Any bidder n who has ever rejected a price is a loser, keeping her station and receiving no payment. Any remaining bidder n who has never rejected a price is a winner, selling her station at price $p_n(h^T)$. Any two functions $p, \hat{p}$ describe different auctions if they have different ending rounds $T \neq \hat{T}$ or specify different prices $p_n(h^t) \neq \hat{p}_n(h^t)$ for some station n that has not exited at some for $t \leq \min(T, \hat{T})$.

Descending clock auctions all share some useful properties.

1. Every descending clock auction is *strategy-proof* and furthermore, also *obviously strategy-proof – a concept defined by my former doctoral student Shengwu Li (2017).* It is "obvious" in his sense that a bidder should accept prices greater than her value v_n and reject all lower offers. This is an "obvious" choice in every situation. For no matter

what price $p_n(h^t)$ is currently offered, if she follows the recommended strategy, her payoff will be at least v_n. If she rejects some price $p_n(h^t) \geq v_n$ or accepts some price $p_n(h^t) \leq v_n$, then her payoff will be at most v_n.

2. Every descending clock auction is *weakly group strategy-proof*. This means that there is no joint deviation by any coalition of bidders that makes every member strictly better off. To see why, consider the first member to deviate. Because the auction is obviously strategy-proof, that bidder cannot strictly benefit from the deviation.
3. Every clock auction p *accommodates budget constraints*, meaning that there is another descending clock auction $\hat{p}$ that leads to total prices of at most any given budget $B \geq 0$ and leads to the same outcome whenever the original auction would satisfy the budget constraint. If the original clock auction is p with termination round T, then such an alternative auction sets $\hat{p}(h^t) = p(h^t)$ for any $t < T$, $\hat{p}(h^T) = p(h^T)$ if the budget constraint is satisfied at T, and, for example, $\hat{p}(h^T) = 0$ if the budget constraint is not satisfied at T.
4. Every clock auction p *accommodates computation time limits*, meaning that there exists another descending clock auction $\hat{p}$ such that $\hat{p}_n(h^t) = p_n(h^t)$ when that can be computed in the allowed time and such that if p always leads to feasible solutions after every history of feasible solutions, then $\hat{p}$ does so as well. This works by the construction that, when computation time expires, set $\hat{p}_n(h^t) = p_n(h^{t-1})$.
5. Every clock auction p *preserves winner privacy*, meaning that if bidder n with value v_n wins at some price of $p_n(h^T)$ when other bidders' values are v_{-n}, then that bidder still accepts all offers and wins at precisely the same price when its value is $\hat{v}_n < v_n$.

In addition to those properties, we showed that the class of clock auctions includes some that approximately minimize the auctioneer's procurement costs and approximately minimize the value of the stations removed from broadcasting, subject to the constraint that sufficiently many stations must be cleared to make the procurement possible.

The computational aspect of this analysis was twofold. The most impressive part was the development of customized algorithms for the FCC problem by Kevin Leyton-Brown and his students at the University of British Columbia. Their algorithms for the checking problem described above ran about 1,000 times faster than the best existing commercial algorithms.

In Leyton-Brown, Newman, Milgrom and Segal (2020), we studied the performance of the FCC reverse auction design by using simulations, with bidder values inferred from the bids in the actual incentive auction. The actual design included some controversial features intended to improve efficiency and reduce costs, in the hopes of clearing more spectrum and raising more net revenue for the Treasury. We evaluated performance criteria including the number of channels cleared, costs of procurement, efficiency of the final allocation, and computation time, and considered how various aspects of the design affected these measures. The summary, which readers should verify by studying our paper, is that the actual FCC design led to high levels of efficiency and expected payments to broadcasters that are much lower than what the FCC would expect using a Vickrey auction.

Conclusion

Auction theory has changed substantially since I made my first studies in what were still its early days. Although the "unrealistic" models of those times have proved their worth in guiding practical auction designs, some of that guidance was off point. In my own work, this showed up in the traditional analysis of the exposure problem. Despite the theoretical worst-case conclusion that exposure problems are intractable, we found that they could sometimes be quite manageable in practice.

For the future, simulations and computational methods are likely to be increasingly important. Still, it still takes theory to understand problems and the scope of proposed solutions. The time has come for old methods and new to work hand in hand.

REFERENCES

Arnosti, N., Beck, M., and Milgrom, P. (2016). Adverse selection and auction design for internet display advertising. *American Economic Review*, **106**(10), 2852–66.

Ausubel, L. M., and Milgrom, P. (2006). The lovely but lonely Vickrey auction. *Combinatorial Auctions*, **17**, 22–26.

Ausubel, L. M., Cramton, P., and Milgrom, P. (2006). The clock-proxy auction: A practical combinatorial auction design. *Handbook of Spectrum Auction Design*, 120–140.

Bichler, M., Milgrom, P., and Schwarz, G. (2020). Taming the Communication and Computation Complexity of Combinatorial Auctions: The FUEL Bid Language.

Bulow, J., Levin, J., and Milgrom, P. (2017). Winning play in spectrum auctions. *Handbook of Spectrum Auction Design*. Cambridge University Press.

Coase, R. H. (1959). The Federal Communications Commission. *The Journal of Law and Economics*, **2**, 1–40.

Day, R., and Milgrom, P. (2008). Core-selecting package auctions. *International Journal of Game Theory*, **36**(3–4), 393–407.

Day, B., and Milgrom, P. (2013). Optimal incentives in core-selecting auctions. *Handbook of Market Design*. OUP Oxford.

Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *The Journal of Finance*, **25**(2), 383–417.

Gale, D., and Shapley, L. S. (1962). College admissions and the stability of marriage. *The American Mathematical Monthly*, **69**(1), 9–15.

Glosten, L. R., and Milgrom, P. R. (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, **14**(1), 71–100.

Grossman, S. J., and Stiglitz, J. E. (1980). On the impossibility of informationally efficient markets. *The American Economic Review*, **70**(3), 393–408.

Hatfield, J. W., and Milgrom, P. R. (2005). Matching with contracts. *American Economic Review*, **95**(4), 913–935.

Kagel, J. H., Lien, Y., and Milgrom, P. (2010). Ascending prices and package bidding: A theoretical and experimental analysis. *American Economic Journal: Microeconomics*, **2**(3), 160–85.

Kagel, J. H., Lien, Y., and Milgrom, P. (2014). Ascending prices and package bidding: Further experimental analysis. *Games and Economic Behavior*, **85**, 210–231.

Kelso Jr, A. S., and Crawford, V. P. (1982). Job matching, coalition formation, and gross substitutes. *Econometrica*, 1483–1504.

Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica*, 1315–1335.

Leyton-Brown, K., Milgrom, P., and Segal, I. (2017). Economics and computer science of a radio spectrum reallocation. *Proceedings of the National Academy of Sciences*, **114**(28), 7202–7209.

Li, Shengwu (2017). Obviously strategy-proof mechanisms. *American Economic Review*, **107**(11), 3257–3287.

Newman, N., Leyton-Brown, K., Milgrom, P., and Segal, I. (2020). Incentive Auction Design Alternatives: A Simulation Study.

Milgrom, P. R. (1979). A convergence theorem for competitive bidding with differential information. *Econometrica*, 679–688.

Milgrom, P. (2000). Putting auction theory to work: The simultaneous ascending auction. *Journal of Political Economy*, **108**(2), 245–272.

Milgrom, P. (2004). *Putting Auction Theory to Work*. Cambridge University Press.

Milgrom, P. (2009). Assignment messages and exchanges. *American Economic Journal: Microeconomics*, **1**(2), 95–113.

Milgrom, P and Eilat, A. (2011). The CAF auction: Design proposal. *WC Docket*, (10–90).

Milgrom, P. and Segal, I. (2002). Envelope theorems for arbitrary choice sets. *Econometrica*, **70**(2), 583–601.

Milgrom, P. and Segal, I. (2020). Clock auctions and radio spectrum reallocation. *Journal of Political Economy*, **128**(1), 1–31.

Milgrom, P., and Stokey, N. (1982). Information, trade and common knowledge. *Journal of Economic Theory*, **26**(1), 17–27.

Milgrom, P. R., and Weber, R. J. (1982). A theory of auctions and competitive bidding. *Econometrica*, 1089–1122.

Myerson, R. B. (1981). Optimal auction design. *Mathematics of operations research*, **6**(1), 58–73.

Ortega-Reichert, A. (1968). Models for competitive bidding under uncertainty.

Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, **16**(1), 8–37.

Vickrey, W. (1962). Auctions and bidding games. *Recent advances in game theory*, **29**, 15–27.

Wilson, R. (1977). A bidding model of perfect competition. *The Review of Economic Studies*, **44**(3), 511–518.

Wilson, R. (1979). Auctions of shares. *The Quarterly Journal of Economics*, 675–689.