Auctions from a Market Design Perspective.

Paul Milgrom, Stanford University, USA

This class treats auction theory emphasizing the aspects of the theory that have proved most valuable in actual market design. There will be seven hours of lectures, with the following preliminary outline.

First hour: Modern Issues in Auction Design.
Based on Milgrom's Nemmers lecture, paper in progress

Hours 2-3: Topics in Classic Auction theory.
Vickrey (1961, 1962). Myerson (1981). Milgrom-Weber (1982, 1982a), Bulow-Roberts (1989), DeMarzo-Kremer-Skrzypacz (2005).

Hours 4-5: The Relationship between Auctions and Matching.
Kelso-Crawford (1982), Demange-Gale-Sotomayor (1986), Shapley-Shubik (1972), Milgrom (2009).

Hours 6-7: Auction Designs for Internet Advertising.
Edelman-Ostrovsky-Schwartz (2007), Varian (2007), Milgrom (2009a).

References:

Athey, Susan, and Glenn Ellison. "Position Auction with Consumer Search." 2007.
Bulow, Jeremy, Jonathan Levin, and Paul Milgrom. "Winning Play in Spectrum Auctions." Stanford University, 2009.
Bulow, Jeremy, and John Roberts. "The Simple Economics of Optimal Auctions." Journal of Political Economy 97, no. 5 (1989): 1060-90.
Demange, Gabrielle, David Gale, and Marilda Sotomayor. "Multi-Item Auctions." Journal of Political Economy 94, no. 4 (1986): 863-72.
DeMarzo, Peter, Ilan Kremer, and Andrzej Skrzypacz. "Bidding with Securities: Auctions and Security Design." American Economic Review 95 (2005): 936-59.
Edelman, Benjamin, Michael Ostrovsky, and Michael Schwartz. "Internet Advertising and the Generalized Second Price Auction: Selling Billions of Dollars Worth of Keywords." American Economic Review 97, no. 1 (2007).
Kelso, Alexander, and Vincent Crawford. "Job Matching, Coalition Formation, and Gross Substitutes." Econometrica 50, no. 6 (1982): 1483-504.
Milgrom, Paul. "Assignment Messages and Exchanges." AEJ Micro 1, no. 2 (2009): 95-113.
———. Putting Auction Theory to Work. Cambridge: Cambridge University Press, 2004.
———. "Simplified Mechanisms with an Application to Sponsored Search Auctions." Game and Economic Behavior Forthcoming (2009a).
Milgrom, Paul, and Robert J. Weber. "A Theory of Auctions and Competitive Bidding." Econometrica 50 (1982): 463-83.
———. "The Value of Information in a Sealed-Bid Auction." Journal of Mathematical Economics 10, no. 1 (1982a): 105-14.
Myerson, Roger B. "Optimal Auction Design." Mathematics of Operations Research 6, no. 1 (1981): 58-73.

Shapley, Lloyd, and Martin Shubik. "The Assignment Game I: The Core." International Journal of Game Theory 1, no. 1 (1972): 111-30.
Varian, Hal R. "Position Auctions." International Journal of Industrial Organization 25 (2007): 1163-78.
Vickrey, William. "Auction and Bidding Games." Recent Advances to Game Theory (1962): 15-27.
———. "Counterspeculation, Auctions, and Competitive Sealed Tenders." Journal of Finance XVI (1961): 8-37.

The Problems and Promise of Auction Market Design

By Paul Milgrom[1]
Department of Economics
Stanford University

Abstract. The years since 1994 have witnessed the emergence of *market design* as a new discipline within economics, in which research and practice exert powerful mutual influences. The problem of designing well-functioning markets has led market designers to revisit fundamental issues such as ways commodities are defined, the ways participants communicate with markets, the trade-off between the type of incentives provided and other attributes of mechanism performance, and the determinants of the scope of markets and whether and how sales of different goods are linked.

Introduction

For economic researchers hoping to influence the rules of real markets, 1994 was an auspicious year. It was a year of redesign for the National Resident Matching Program, whose rules were revised based on the recommendations of an academic researcher (Alvin Roth). And, it was the year of the first US radio spectrum auctions, the rules for which were provided in proposals by economic researchers (especially one by Paul Milgrom and Robert Wilson and a similar one by Preston McAfee) who also later advised bidders in the auctions. Both of these designs proved highly influential, inspiring similar designs in other jurisdictions. The exchanges between researchers and others who collaborated on those designs foreshadowed the active engagement of research economists during the following years in setting up or improving other markets as well, and in learning from those about gaps in our understanding of the details of *market design*.

In the decade or so before 1994, the similarly named field of *mechanism design* had already emerged as an important one in economic theory and game theory. It emphasized that the rules of auctions, political institutions, tax and regulatory systems, and much more can be objects of choice, that the possible outcomes of different mechanisms can be conceived as limited by game-theoretic equilibrium concepts ("incentive constraints"), and that different rules can be evaluated based on the mapping they promote from environments to outcomes. Auction theory became a poster child for mechanism design theory, because the rules of auctions are typically very precise and easily modeled and are often purposefully chosen by

[1] Parts of this lecture were first delivered as part of the WEAI Presidential Lecture in July, 2008. This revised version was delivered as the Nemmers Prize Lecture at Northwestern University in October 2009.

the auctioneer.[2] Yet even a brief attempt to create real auction mechanisms makes it apparent that mechanism design theory had been silent about many of the questions that arise in practice. As Roth (2002) has emphasized, market design is a kind of economic engineering that demands immediate answers to practical questions, so silence about even poorly understood design issues is not an option for market designers. Solutions must be invented on the spot, which can examined carefully only later. My plan in this lecture is to revisit four of the key issues I have encountered recurrently in auction market design.

The first issue concerns *product definition*. In the traditional economic models, , the set of goods is treated as given to the modeler before any analysis begins. But in many real-world auction problems, the first questions are: what are we selling (or buying)? How is the commodity to be defined or the terms of the contract set? This is a very common question in practice, but one about which mechanism design theory has been largely silent.

A second issue concerns *messages*. In the usual models of economic theory, a type space is given and, at least in the simplest direct mechanisms, participants merely report their so-called "value types," which describe the values they place on the various possible outcomes. Game theorists have put growing emphasis on the special assumptions this approach entails, but not on the impracticality of reporting complete value information in problems of even modest size. (With just twenty items for sale that can be bought in any combination, a value-type is a vector with more than a million entries!) That impracticality leaves us with the question: if types cannot be completely reported and described, what messages should be used? How does the limit affect what can be achieved?

The third issue concerns the need to provide *incentives*. Mechanism design theory commonly treats incentives as a constraint on the mechanism design, but that is best regarded as a mere analytic device. We don't know precisely how different kinds of incentives affect participant behavior, and so good market design involves exploring the trade-offs between incentive-compatibility of various kinds – including notions of approximate incentive compatibility – and other design goals. How much can be gained by relaxing exact incentive constraints to approximate ones?

A final issue is about *linkages among markets*. Economic theory research takes two approaches to about this. Mechanism design research typically adopt a small-worlds view, focusing on just one decision or the allocation of a single resource, without asking how that decision came to be separated from all the potentially related ones. General equilibrium theory stands at the opposite pole, treating all decisions for all products and all time as interlinked in one grand optimization. In real markets, a buyer or seller must decide: what should be the scope of this auction? For example, should an electric utility or system operator buy or sell

[2] Matching – particularly of students to schools – became the other great area of applied mechanism design research, for a similar reason.

power, peaking capacity and transmission capacity in one auction, or in several separate ones?

Separate sections of this paper discuss each of these topics.

Product Definition and Conflation

One simple fact about physical commodities is that no category is ever actually homogeneous. The most common view among economists, articulated by Debreu, (1959), is that commodities are distinguished not only by their physical characteristics but also by the time and place of their availability (and sometimes also by the contingency under which they will be delivered). In that formulation, it is particularly clear that no category can include more than a single item: no two items can be available at precisely the same time and place![3]

Electric power in California provides a good example. One characteristic that distinguishes power is its time of availability. On a hot day in Los Angeles, extra power generated at night cannot run the air conditioners when they are needed in the afternoon. Location matters, too. Unless there is sufficient transmission capacity, which is costly to build, maintain and operate, power generated in northern California or out-of-state is useless for running air conditioners in southern California. And with outages potentially caused by storm, earthquake, flood, local equipment problems, etc, there are no two locations for which power is actually the same, and contingencies matter, too.

Does this mean that the market must recognize all these distinctions for pricing power, with a separate price for every combination of time, place and contingency? Clearly not. As for many commodities, electricity markets are designed incorporating general principles of *standardization* and *conflation* – treating unlike things as if they were the same – and *deconflation* – recognizing more distinctions for some purposes than are incorporated into the standard. For electricity, the system may conflate power from different sources and times by applying a single price to all of it, and the system operator may also deconflate, recognizing many more distinctions for both planning and decision-making.

One advantage of conflation in electricity markets is that it reduces the importance of indivisibility for the market, making it more effective to use prices to guide resource allocation. The corresponding theory is associated with the Shapley-Folkman theorem, which holds roughly that the “size” of the non-convexities in the social production set increase with the number of goods in the economy.[4]

Another helpful example of conflation is the market for wheat. Wheat is a textbook example of a commodity, and many of the once-important distinctions made in wheat markets are easily forgotten or minimized. Yet before wheat markets

---

[3] Debreu was probably aware of this difficulty, writing: “A commodity is characterized by its physical properties, the date at which it will be available, and the location at which it will be available.” By using “date” instead of “time,” he formally sidesteps the simultaneity problem and allows multiple copies of each commodity, but this formulation raises other issues, as the main text explains.
[4] The literature about the Shapley-Folkman theorem begins with Starr (1969). For the relationship between the number of commodities and the duality gap, see Aubin (1976) and Cassels (1975).

evolved modern standards, buyers regarded the product of each farm, and even of fields within each farm, as heterogeneous. Describing the behavior of buyers in the spot market for wheat of the late nineteenth century, Peter Dondlinger writes, "[F]or each transaction they would analyze a sample to determine its value. The measurement costs were very high."[5]

Without standardization wheat markets were cumbersome and expensive and futures contracts were difficult or impossible to write. Many of the important aspects of the future delivery (say for "clean, dry, sound wheat") are highly subjective. Until standards emerged (partly under pressure from the powerful railroads), participation in wheat markets was expensive, local, and fragmented.

The emergence of standards did not eliminate all heterogeneity among farmers and their crops. If a standard required that a bushel of wheat weighed at least X pounds and the fraction of impurities was at most Y%, then buyers would still pay more for standard wheat that weighed more or had fewer impurities and might pay a premium for it. But a system of inspection could still emerge – and did emerge – that allowed many low cost transactions and enabled the development of futures contracts.

To many readers, the market for uncut diamonds will seem to be the polar opposite of that for wheat. They think of each diamond as being physically unique, varying in size, shape, color, imperfections, and so on, and so needing to be priced individually. And, indeed, individual diamonds do need to be evaluated physically to determine the most valuable cuts. Nevertheless, just as for electricity and wheat, buyers recognize the possibility of substituting one diamond for another. How can these inspection costs be mitigated?

BHP Billiton uses an auction to sell diamonds, allowing buyers to inspect only one standard *split* in each category ("deal") and using a uniform price auction to sell splits. That is conflation. But this market also turns principles of standardization on their head to apply deconflation in its pricing. Billiton adjusts the uniform price in each deal by applying its own relative value index. It compares each split to the standard and multiplies the price by the corresponding ratio to set a individual price for each split.

The importance of conflation for pricing is found even in financial markets. Bengt Holmstrom (2010) has argued that what makes AAA-rated securities so valuable is that buyers do not feel the need to evaluate them individually. High-rated securities are supposed to have no significant potential for adverse selection, so buyers can purchase without inspection and markets can be liquid and thick. After the financial crisis of 2008-09, securities that were once thought to be safe were found to be varied and risky, destroying conflation, increasing the need for costly inspection, and drastically reducing the volume of trade.

See Levin and Milgrom (2010) for a discussion of the possibilities for conflation and deconflation in the market for Internet advertising.

---

[5] Dondlinger (1908).

The topic of conflation and standardization is pervasive in market analysis, and comes up repeatedly in the following sections.

Message Spaces and Simplicity

One of the most celebrated results of mechanism design theory is the *revelation principle*, which is often interpreted to say that any performance that can be achieved in equilibrium by any mechanism can be achieved by an *incentive-compatible direct mechanism*, which is a mechanism in which participants are simply asked to report everything they know or believe, the outcome is computed from those reports, and the participants' incentives lead them to report truthfully. This remarkable principle has allowed researchers to write mechanism design problems in a simple and tractable form, paving the way for useful analysis.

The revelation principle seems to suggest that practical mechanisms can be designed just by asking each participant to make a one-time report listing its values for all possible outcomes and selecting the outcome based on the collected reports. But in fact, even when participants' only relevant information is just their values for the different goods assignments, the reporting requirements of a direct mechanism can be impossibly burdensome. For example, in one of the large FCC radio spectrum auctions (auction #66) with L = 1132 licenses for sale and N > 100 bidders participating, each bidder can receive any of 2^L–1 possible non-empty assignments for itself and, if the bidder also cares about the joint goods assignment, the number of possible outcomes is $(N+1)^L$. These are impossibly large numbers, and we can dismiss without further analysis any possibility that a bidder might report (or even know) the values for so many different assignments.

What sorts of direct reporting mechanisms are practical? There are two pure approaches. The first is to employ a mechanism like an ascending auction, in which the auctioneer elicits information in sequential fashion over time. Theoretically, this approach can sometimes economize on information by seeking reports only about values in relevant ranges (Parkes, 2005 and Compte and Jehiel 2000). To the extent that participants need to report their values only for relevant outcomes, such mechanisms also economize on attention, communication, and computing power. Yet the auctioneer's decision about what information to seek can be a difficult one, and the practical problem of deciding what feedback to give to bidders to guide their decision has received little attention (but see Kwasnica, Ledyard et al. (2005)).

The second pure approach, which is quite common, works by limitin the messages that participants can report, based on *a priori* decisions about what is important. For example, in the National Resident Matching Program that matches doctors to hospitals in the United States, hospitals may have preferences over the composition of the class of doctors that they hire, but can make only limited reports characterized mainly by rank-order list of candidates and the number of openings at the hospital.

The idea of limiting messages to simplify reporting is found sometimes in auctions as well. For example, in bidding to place ads on a search results page corresponding to a particular search phrase, bidders must typically name a price per click to be paid each time and searcher visits the bidder's site. The advertiser typically cannot specify different prices for different positions on the page, and often cannot make the price it offers contingent on demographic characteristics of the searcher. The limited messages used for these auctions remind us of the important role played by conflation in real markets, where distinctions are intentionally suppressed to thicken markets, reduce adverse selection, and so on.

But if conflation and other forms of simplification are sometimes beneficial, how do they affect the equilibrium outcomes of the mechanism? Clearly, simplification can eliminate equilibria by eliminating strategies that are played at equilibrium. But it can have the reverse effect, too, adding equilibrium profiles by eliminating the profitable deviations. If the original mechanism is a good one, then adding equilibrium profiles can be dangerous, by adding the possibility of undesirable outcomes. Can we avoid that possibility when designing simplified mechanisms? Can we do that even when the simplification is severe?

The trade-offs appear to be tricky. On one hand, limiting a participant's messages is more likely to eliminate a profitable reply when the simplification is severe. On the other, severely limiting competitors' when competitors' messages correspondingly reduces a player's need for strategic flexibility. Can the two effects be balanced? A precise answer to the question hinges on the *outcome closure property*, for which we first need to introduce some definitions and notation.

We begin with a certain mechanism that we might like to simplify. Suppose that there are N participants who make reports in the sets $M_1,\ldots,M_N$. A *report profile* is a point in $M= M_1\times\ldots\times M_N$. For each participant n, the relevant part of the outcome is a point in a set X_n, and the outcome is a point in $X\subseteq X_1\times\ldots\times X_N$. A *mechanism* is a pair (M,ω) where $\omega:M\rightarrow X$ determines the outcome as a function of the report profile.

When preferences/payoffs are specified, the mechanism formally becomes a game. We associate payoffs with outcomes by allowing each participant n to have an outcome-dependent payoff function $u_n:X_n\rightarrow\Re$. This formulation is quite general and subsumes both public decision problems and private allocation problems.

A *simplification* of mechanism (M,ω) is another mechanism $(M',\omega_{|M'})$ in which the set of messages available to each participant is reduced: $M'\subseteq M$. We may also say that (M,ω) is an *extension* of $(M',\omega_{|M'})$. Generally, when we observe what we believe to be a simplification, we interpret its characteristics by reference to some extension, even though extensions, even with a particular message space, are never unique.

Informally, the condition we will need is that when j's competitors are restricted by a simplification, then for any simplified message profile they report, any outcome that participant j could bring about from any message in M_j can be achieved, or at least closely approximated, by using just a simplified message. To state this condition precisely, we need to assume that the outcome space has some topology. Then, a simplification has the *outcome closure property* if for every player *j* and every profile of restricted messages $m_{-j}\in M_{-j}$, $\text{cl}(\omega(M_j,m_{-j})) = \text{cl}(\omega(M'_j,m_{-j}))$, where "cl" denotes the closure of the set.

Simplification Theorem.[6] Let $(M',\omega_{|M'})$ be a simplification of $(M',\omega_{|M'})$ satisfying the outcome closure property. Let *u* be any profile of continuous utility functions and let $\varepsilon \geq 0$. Then every report profile *m* that is a (full-information) ε-Nash equilibrium profile of the simplified mechanism $(M',\omega_{|M'})$ is also a full-information ε-Nash equilibrium of the extended mechanism (M,ω). Conversely, if $(M',\omega_{|M'})$ be a simplification of $(M',\omega_{|M'})$ *not* satisfying the outcome closure property, then there is a profile *u* of continuous utility functions and a report profile *m* that is a (full-information) Nash equilibrium profile of the simplified mechanism $(M',\omega_{|M'})$ but not of the extended mechanism (M,ω).

As a first example, consider the National Resident Matching Program, in which hospitals can report only rank-order lists of doctors and a number of openings, while doctors can report just a rank-order list of hospitals. The outcome function then selects a stable allocation for the reported preferences. Suppose we consider the extended mechanism in which hospitals can report any preference ranking according to which doctors are substitutes. This is known to be the largest message space that guarantees the existence of a stable allocation. With a bit of thought, it is clear that if other hospitals are limited to reporting rank order lists, then any group of doctors that a hospital could hire by any extended report can also be hired by listing the same doctors in any order as the only acceptable ones in a simplified report. So, the outcome closure property is satisfied and we can conclude that the simplification does not introduce any additional equilibria (and, in particular, undesirable ones).

A similar argument can be made for pay-as-bid combinatorial auctions. If my competitors must name individual prices for each item they want in an auction and if I have some bid that can acquire a particular set S at a total price p, then p must be more than the sum of my competitors' highest bids for the elements of S. So, I can divide p among the items to win each one in the set S at a total price of p: the outcome closure property is satisfied.

These are severe simplifications, so it is perhaps surprising that they do not introduce any new Nash equilibria or any new approximate equilibria.

The use of simplification can also lead to positive improvements, essentially by eliminating bad equilibria that can be all too likely when reporting is complex. We illustrate this possibility by summarizing a model of online search advertising of (Milgrom, 2009), in which there are several ad positions offered on a search page, decreasing in value from top to bottom. Each buyer is limited to purchase just one.

In the model, bidders may have different values for clicks originating from different ad positions. Each bidder incurs a small positive cost for each price reported in its message. For the simple mechanism, a bidder reports a single price-per-click which applies to all advertising positions on the search page. For the extended mechanism, the bidder reports a price-per-click separately for each position. The mechanism conducts a series of second-price auctions among the

[6] The direct part is due to Milgrom (2009), the converse part to Perez (2009).

remaining bidders, beginning with the top position on the page. This is the mechanism described by Edelman, et al (2007) and Varian (2007).

Because bidding is costly, in any full information pure Nash equilibrium, no losing bids are made. According to that solution concept, the extended mechanism raises zero revenue. The simplified mechanism fails to attract a bid from the $N+1^{th}$ bidder needed to generate fully competitive prices, but always generates positive revenue and, when there are just N serious bidders, generates the competitive revenue at a pure strategy equilibrium. In contrast, with N positions on the page, the simplified mechanism raises a competitive revenue from a market with N bidders.

The general point that this simple model illustrates is that thickness is important for pricing. Conflation reduces the number of markets and increases the number of transactions in each. When participation is costly, unconflated markets are inevitably thin in ways that impairs their operations.

Incentives for Participants

Another central consideration for mechanism design is the provision of incentives, but the best way to account for incentives is extremely subtle. On one hand is the debate about game theoretic solution concepts as they apply to very rational agents. On the other is the fact that participants' response to incentives is much murkier in practice than in any of the usual models of economic theory. This is hugely important, because there can be severe trade-offs between the nature of incentives provided and the mechanism's other attributes.[7]

Much of the research in mechanism design restricts attention to participants with quasi-linear utility, meaning that each participant's payoff is the value of the decision or goods assignment plus-or-minus the value of any cash transfer the participant receives or makes, and the review here generally takes that assumption as given.

To illustrate some of the restrictions implied by incentive requirements for other mechanism attributes, we begin by reviewing what can be achieved by a mechanism designed to ensure *ex post* incentive compatibility. This incentive concept means that if any participant could learn the other participants' information and if the messages reported by others are truthful, then the participant would have no incentive to falsify his own report. This concept would appear to be a very attractive one, because it eliminates not only the incentives to misreport values, but also participants' incentives to spy on others or to waste resources trying to learn or guess about what others know. Information of that sort is simply irrelevant to the participant's decision.

While the concept appears potentially attractive, there is a large set of environments for which it is too much to ask: it eliminates all but the most trivial mechanisms. More specifically, (Jehiel, Meyer-ter-Vehn et al. 2006) have proved

---

[7] Sometimes, it is the absence of trade-offs that is most surprising. Myerson (1981) shows that the maximum expected revenue from an auction in certain environments is the same regardless of whether the bidders are to be provided with dominant strategy incentives or Bayesian incentives.

that, apart from exceptional (non-generic) environments, if participants have multi-dimensional information and interdependent values, then any *ex post incentive* compatible mechanism has a constant outcome – one that is entirely independent of the participants' information. This is an impossibility theorem: it says that nothing of much interest can be achieved by such a mechanism. Mechanisms with this property waste any valuable information that the participants may have.

The problem in these cases is not that *ex post* incentive compatibility is wrong or uninteresting. If the environment allows us to construct mechanisms with such clear incentives and other good properties, we might like to do so. The "non-generic" case of private values would seem to offer an interesting opportunity. If we combine the requirements that, over private values environments, the mechanism is *ex post* incentive compatible ("straightforward"), allocates goods to maximize the reported value of the outcome, and charges a zero price to those who acquire no goods, then according to theorems of Green and Laffont (1977) and Holmstrom (1979), there is just one mechanism that satisfies all the requirements: the (generalized) Vickrey auction.[8]

The Vickrey, too, has properties that make it unsuitable for many applications. Ausubel and Milgrom (2006) point out several of the most damaging one. We illustrate a few of these with variations of the following simple example.

Table 1			
Bidders	Item A	Item B	Pair AB
1	0	0	10
2	10*	9.99	10
3	9.99	10*	10

In this example, the unique goods allocation that maximizes total value assigns A to bidder 2 and B to bidder 3, creating a total value of 20. It is routine to verify that the Vickrey formula specifies that each bidder pays zero. This is the *low revenue problem*. Since all three bidders were each willing to pay 10 for the pair AB, there is ample competition for the pair, so zero revenue is a serious problem. Put another way, the Vickrey allocation would seem to be *blocked* in any real market by a coalition of the seller and bidder 1, who could do a deal that is better for both of them. One possible goal of mechanism design is to seek outcomes that are unblocked (and hence in the core of the corresponding trading game).

Related is the *shill bids problem*. Table 1 could be the result of a deception. Perhaps there are really just two bidders, 1 and 2, but bidder 2 has been clever, inventing a third bidder and submitting reports under two identities. The situation might be one in which bidder 2 should win the items, but prices should be higher. Or perhaps bidder 2 should not win at all. Its values might be much lower than shown,

[8] The theorems are usually formulated in terms of public goods problems, in which case the mechanism might is the Vickrey-Clarke-Groves mechanism after Vickrey (1961), Clarke (1971), and Groves (1973).

but the reports in Table 1 allow it to acquire both goods at a total price of zero. "The game is bigger than you think" is good slogan for market design: participants are sometimes quite clever in finding ways to cheat or disrupt a mechanism.

Yet another issue raised by Table 1 is the *problem of disqualifications*. In many auctions, bidder qualifications are checked fully only for winners and only after the results are in. In Table 1, the auctioneer would benefit financially by disqualifying bidder 3, thereby raising its revenue from zero to 10. No commonly used auction mechanism provides quite the same incentive.

Table 2			
Bidders	Item A	Item B	Pair AB
1	0	0	10*
2	4	4	4
3	4	4	4

The *problem of collusion* is illustrated by comparing tables 1 and 2. Suppose that Table 2 represents the real values of the bidders. With honest reporting, bidder 1 is the winner and the Vickrey price is 8. But, in a Vickrey mechanism, the losing bidders can collude profitably, raising their bids to match the numbers in Table 1. In that way, both become winners at zero prices.

A similar example of too-strict incentives limiting what can be achieved is found in matching, rather than in auctions. For the course allocation problem, in which students have preferences over sets of courses, the only straightforward mechanisms are versions of a serial dictatorship, in which some student chooses all of her courses first, a second student chooses all of his from the remaining openings, and so on, to the final student who chooses all of her courses last, from whatever remains. Such a mechanism can lead to outcomes that seem very unfair, as for example when one student gets all the popular courses and another student gets none, but that sort of outcome is unavoidable if we limit attention to *ex post* incentive-compatible mechanisms.

Practical mechanism design sometimes seeks to use weaker notions of incentive-compatibility, including notions of approximate incentives, to overcome these limitations.

One practical example is the Day-Milgrom pricing rule (Day and Milgrom, 2007), versions of which have been incorporated into spectrum auction designs in the UK, Ireland, Portugal, Denmark and the Netherlands. The rule applies to mechanisms in which value reports about some goods (such as licenses to use slices of the radio spectrum) are collected either by a direct mechanism or in a series of rounds. Like the Vickrey mechanism, the Day-Milgrom core-selecting auctions assign the goods to maximize the total reported value of the allocation and prices are computed roughly as follows. The auctioneer uses the reported values to specify a trading game, identifies the core of that game, and identifies a point that minimizes the seller's total revenue subject to being in the core. When there are multiple such points, a tie-break rule is applied. This pricing rule that provides the smallest total gain to false reporting by bidders over all core allocations, starting from the point of truthful reporting.

The Day-Milgrom price rule is closely connected to the Vickrey rule. When the Vickrey outcome for any set of bids is the core, it is provably the unique revenue-minimizing core point, so the Vickrey and Day-Milgrom prices are then identical. In such cases, truthful reporting is an *ex post* equilibrium of the corresponding mechanism. This is always the case when goods are substitutes for the bidders (Ausubel and Milgrom, 2005).

But in examples like that of Table 1, for which the Vickrey revenue of zero is too low to be part of a core allocation, the Day-Milgrom rule forces the total price to be higher (equal to 10), an outcome that sacrifices *ex post* incentives in favor of revenue for the seller. If bidders are subject to a great deal of uncertainty about others' values, the incentive to deviate substantially from truthful reporting may be small, because the maximum gain from a clever deviation has been minimized (subject to the core) and may be zero, and identifying the profitable deviation may be hard.

At least since the pioneering work of Roberts and Postlewaite (1976), there has been a tradition of analyzing market mechanisms in terms of approximate incentive-compatibility. The idea is that in environments with many participants, no single bidder can have much influence on prices, so a bidder cannot do much better than to report truthfully and let the mechanism assign her the best package of goods at the final market prices.

Recent work applies the approximate incentives idea to matching markets with or without the explicit use of prices. Budish (2009) is perhaps most explicit, devising an approximate price-mechanism for course matching starting with nearly equal endowments and showing that, in large markets, participants cannot much affect prices and so can gain only little by reporting falsely. Recall that this is the very problem for which *ex post* incentive compatibility leads to the extreme outcomes of serial dictatorships. The contrast is striking.

Kojima and Pathak (2008) and Kojima and Minea (2009) find similarly that incentives for false reporting are minimal in matching markets without explicit prices. In examples, drawn from both auctions and matching, there can be a very valuable payoff to replacing exact incentives with some weaker or approximate notion of incentives in a market design.

Linkages Among Markets

Our last topic is linkages among markets. Many commodity markets in the world are pure: there is no direct linkage between the price setting for one commodity and another. Yet that seems sharply at odds with the thinking that underlies consumer theory and general equilibrium theory. In those theories, demand in each market depends on all prices. How can each be set separately?

Thinking about linkage leads immediately back to the thoughts and insights of the previous sections.

First, for physical goods, no two items can ever be exactly the same in physical characteristics, time of availability, and place of availability. All such markets involve standards and conflation, so all involve creating some linkages, at least among items that are similar. For sufficiently similar items, that might be done by pure conflation (suppressing and ignoring differences), conflation and deconflation (as in the BHP Billiton diamond auction), or by linking the sale of related products, especially substitute products.

One example is the export market for powdered milk from New Zealand. The monthly auctions allow bidders to bid, in a simultaneous ascending clock auction,[9] for various types and contracts for milk powder. The powder can be instantized or not, treat at ultra-high temperature (UHT milk) or not, sold for delivery next month or at a different time. It can originate from New Zealand or Australia. Bidders can choose according to their preferences and the relative product prices at each round of the auction.

The simultaneous multiple round auction (SMRA) was one of the 1994 innovations that helped launch the field of market design. Simultaneous, multi-product auctions had been to the FCC by Milgrom and Wilson and by McAfee, but the regulators had feared that the design might work poorly if bidders, eager to let others set prices before they committed themselves to bid, withheld their bids until late in the auction. The design innovation that solved that problem was Milgrom's *activity rule*, which required bidders to be active early in the bidding to maintain eligibility to bid later. Versions of the activity rule are now widely used in these auctions around the world.

Radio spectrum auctions are another example of ones in which related but different items are sold, but in this case the items are not necessarily substitutes. A bidder may want to acquire a particular 10MHz spectrum license only if it can acquire an adjacent license as well. Or, it might want coverage in a particular city only if it can acquire coverage in nearby cities as well. The licenses in that case are complements.

A second issue with linked markets concerns the message space to be used in the auction. With multiple items, how should bids for packages be described? Can values be accurately described with concise messages? My own research has recently introduced a practical message space – *assignment messages* – for auctions of products that are substitutes. The assignment auction, which in its simplest form aims to replicate and improve the results of an SMRA but with a sealed bid auction, is described in Milgrom (2009).

The third issue with linking markets is incentives. We have already seen how one theoretically attractive auction design (the Vickrey auction) can have

[9] A clock auction is an auction in which the auctioneer calls the price and bidders respond with demands. The *Walrasian tatonnement*, familiar to economic theorists, is a kind of clock auction.

performance problems when goods are complements. As we have already discussed, the Vickrey auction has uniquely good incentive properties for the most commonly studied economic environments. Yet it leads to the problems described in the example of Table 1. This problem is particular to designing markets for goods that are complementary: there is no similar example of a Vickrey outcome with zero revenue (or less than minimum core revenue) when goods are substitutes for all the bidders.

Conclusion

Market design is an engineering discipline linked to mechanism design. The practice of designing markets highlights several questions that are rarely studied in mechanism design. In this lecture, I offer some initial thoughts about several such questions, including how commodities are defined, how participants communicate with markets, what sorts of incentives a mechanism should provide, and why there are linkages among some markets but not among others. All of these questions merit more study.

References

Aubin, J.P. and I. Ekeland (1976). "Estimates of the Duality Gap in Nonconvex Optimization." *Mathematics of Operations Research* **1**(3): 225-245.

Ausubel, Lawrence and Paul Milgrom (2006). The Lovely but Lonely Vickrey Auction. *Combinatorial Auctions*. P. Cramton, Y. Shoham and R. Steinberg. Cambridge, MA, MIT Press.

Budish, Eric (2009). The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes.

Cassels, J.W.S. (1975). "Measures of the Nonconvexity of Sets and the Shapley-Folkman-Starr Theorem." *Mathematical Proceedings of the Cambridge Philosophical Society* **78**(39): 433-436.

Clarke, E.H. (1971). "Multipart Pricing of Public Goods." *Public Choice* **XI**: 17-33.

Compte, Olivier and Philippe Jehiel (2000). "On the Virtues of the Ascending Price Auction: New Insights in the Private Value Setting." *Working Paper, CERAS*.

Day, Robert W. and Paul Milgrom (2007). "Core-Selecting Package Auctions." *International Journal of Game Theory* **36**: 393-407.

Debreu, Gerard (1959). Theory of Value: An Axiomatic Analysis of Economic Equilibrium. New Haven, Yale University Press.

Dondlinger, Peter (1908). The Book of Wheat: An Economic History and Practical Manual of the Wheat Industry. New York, Orange Judd Company.

Green, Jerry and Jean-Jacques Laffont (1977). "Characterization of Satisfactory Mechanisms for the Revelation of Preferences for Public Goods." *Econometrica* **45**: 427-438.

Groves, Theodore (1973). "Incentives in Teams." *Econometrica* **61**: 617-631.

Holmstrom, Bengt (1979). "Groves Schemes on Restricted Domains." *Econometrica* **47**: 1137-1144.

Holmstrom, Bengt (2008). "Discussion of Gary Gorton's 'The Panic of 2007'."

Jehiel, Philippe, Moritz Meyer-ter-Vehn, Benny Moldovanu and William Zame (2006). "The Limits of Ex Post Implementation." *Econometrica* **74**(3): 585-610.

Kojima, Fuhito and Mihai Minea (2009). Incentives in the Probabilistic Serial Mechanism.

Kojima, Fuhito and Parag Pathak (2008). Incentives and Stability in Large Two-Sided Matching Markets.

Kwasnica, Anthony, John O. Ledyard, David P. Porter and Christine DeMartini (2005). "A New and Improved Design for Multi-Object Iterative Auctions." *Management Science* **51**(3): 419-434.

Levin, Jonathan and Paul Milgrom (2010). "Online Advertising: Heterogeneity and Conflation in Market Design." *American Economic Review* **100**(2): 603-607.

Milgrom, Paul (2009). "Assignment Messages and Exchanges." *AEJ Micro* **1**(2): 95-113.

Milgrom, Paul (2009). "Simplified Mechanisms with an Application to Sponsored Search Auctions." *Game and Economic Behavior*.

Myerson, Roger B. (1981). "Optimal Auction Design." *Mathematics of Operations Research* **6**(1): 58-73.

Parkes, David (2005). "Auction Design with Costly Preference Elicitation." *Annals of Mathematics and Artificial Intelligence* **44**(3): 269-302.

Perez-Richet, Eduardo (2010). "A Note on the Tight Simplification of Mechanisms."

Roberts, John and Andrew Postlewaite (1976). "The Incentives for Price-Taking Behavior in Large Exchange Economies." *Econometrica* **44**(1): 115-129.

Roth, Alvin E. (2002). "The Economist as Engineer: Game Theory, Experimental Economics and Computation as Tools of Design Economics." *Econometrica* **70**(4): 1341-1378.

Starr, Ross (1969). "Quasi-Equilibria in Markets with Non-Convex Preferences." *Econometrica* **37**(1): 25-38.

Vickrey, William (1961). "Counterspeculation, Auctions, and Competitive Sealed Tenders." *Journal of Finance* **XVI**: 8-37.