BUSINESS REVIEW

ISSN 0007-7011

Federal Reserve Bank of Philadelphia

MARCH • APRIL 1988

BUSINESS REVIEW

The BUSINESS REVIEW is published by the Department of Research six times a year. It is edited by Judith Farnbach. Artwork is designed and produced by Dianne Hallowell under the direction of Ronald B. Williams. The views expressed herein are not necessarily those of this Reserve Bank or of the Federal Reserve System.

SUBSCRIPTIONS. *Single-copy subscriptions for individuals are available without charge. Institutional subscribers may order up to 5 copies.*

BACK ISSUES. *Back issues are available free of charge, but quantities are limited: educators may order up to 50 copies by submitting requests on institutional letterhead; other orders are limited to 1 copy per request. Microform copies are available for purchase from University Microfilms, 300 N. Zeeb Road, Ann Arbor, Michigan, 48106.*

Please send subscription orders, back orders, changes of address, and requests to reprint to the Federal Reserve Bank of Philadelphia, Department of Research, Publications Desk, Ten Independence Mall, Philadelphia, PA 19106-1574, or telephone (215) 574-6428. Please direct editorial communications to the same address, or telephone (215) 574-3805.

MARCH/APRIL 1988

Investigating How Markets Work: Two Perspectives

Many of us in the United States give little thought to how the prices of the products and financial assets that we buy are determined. When the price of coffee goes up, people say, "it's the law of supply and demand." The concepts are as old as economics itself.

Although the U.S. is often characterized as a "capitalist" country where we have "free markets" and "competition" that allow the forces of supply and demand to set prices for goods and services, many Americans don't stop to consider how, or whether, such mechanisms actually work. We take for granted that markets work—at least that they work well enough so that, when we arrive to buy goods or services, we don't have to stand in line a long time. This is not the case everywhere, however, and some countries, such as the Soviet Union and China, are actually trying to introduce more market forces of supply and demand into their economies.

What is it that makes markets function well? And how do alternative types of price-setting mechanisms perform differently? Economists continue to study these questions, and have been going beyond the introductory textbook explanation of supply and demand. This issue of the *Business Review* is devoted to explaining two approaches used in recent economic research on how markets function. One approach is empirical, but uses experimental methods—like a lab experiment, it sets up simplified markets in a laboratory environment and observes how prices adjust to equate supply and demand. The other approach is theoretical—it models auction markets as strategic games, and finds that the rules of the game affect how supply and demand determine price.

Richard W. Lang
Senior Vice President & Director of Research

Going, Going, Gone: Setting Prices With Auctions

*Loretta J. Mester**

Each week when the U.S. Treasury auctions off billions of dollars of Treasury bills, it is setting prices with a mechanism that is over 2000 years old. Auctions are used to sell a wide range of objects, from art works to drilling rights to government contracts—and the stakes are high. On November 13, 1987, the *Wall Street Journal* reported that Van Gogh's "Irises" was auctioned for a record $53.9 million dollars (beating the March 1987 record of almost $40 million paid for Van Gogh's "Sunflowers"). On September 30, 1980, U.S. oil companies paid $2.8 billion in an auction for drilling rights on 147 tracts in the Gulf of Mexico. In another auction, a consortium of Phillips Petroleum and Chevron USA bid $333.6 million to win the offshore drilling rights near Point Arguello, California; their bid was over twice the next highest offer.[1]

*Loretta J. Mester is a Senior Economist in the Banking and Financial Markets Section of the Research Department of the Federal Reserve Bank of Philadelphia.

[1]The drilling rights examples are from Paul R. Milgrom and Robert J. Weber, "A Theory of Auctions and Competitive Bidding," *Econometrica* 50 (September 1982), and Marc Levinson, "Using Science to Bid for Business," *Business Month* (April 1987), respectively.

The longevity of auctions, the variety of objects sold through auctions, and the enormous volume of business done in auctions today point to the importance of understanding the economics of auctions. Economists have developed theoretical models to try to answer a number of questions about auctions. What are the advantages of using an auction to set the price of an object? Which set of auction rules should the seller use? And given a particular set of auction rules, what is the buyer's best bidding strategy? The answers to these questions will be important to anyone who participates on either side of an auction transaction, including government procurers, government debt managers, investors, collectors, and businesses.

FINDING THE RIGHT PRICE

Auctions as Price-Setting Mechanisms. Auctions, which have been used since ancient times (see 2500 YEARS OF AUCTIONS...AT A GLANCE), are one of the basic mechanisms for determining the prices of goods to be exchanged.[2] Two other mechanisms are posted prices and negotiated prices. Retailers usually post a price for each good they want to sell, and individual buyers have little choice but to take it or leave it. Sellers find posted prices inexpensive to manage, but in the short run, they are inflexible to changes in demand or to changes in an individual buyer's tastes since there is a lag between the time the price is set and the product is sold. Also, posted prices cannot reflect subtle variations in quality among different units of a particular product being sold. When prices are negotiated, both buyers and sellers influence the price substantially, haggling back and forth. Manufacturers usually negotiate the price of

[2]This discussion is based upon Paul R. Milgrom, "The Economics of Competitive Bidding: A Selective Survey," in Leonid Hurwicz, David Schmeidler, and Hugo Sonnenschein, eds., *Social Goals and Social Organization: Essays in Memory of Elisha Pazner* (Cambridge: Cambridge University Press, 1985).

2500 Years Of Auctions...At A Glance

One of the earliest reports of an auction was by Herodotus who described the bidding of men for wives in Babylon around 500 B.C.[a] This auction was unique since bidding sometimes started at a negative price.[b] Some scholars interpret the Biblical story of the sale of Joseph into slavery as an even earlier reference.[c] In ancient Rome, auctions were used in commercial trade and were held in the *atrium auctionarium* where goods could be displayed prior to sale. Auctions were also used to liquidate property by Romans in financial straits. Caligula auctioned off family belongings to cover his debts and Marcus Aurelius held an auction of royal treasures to finance a state deficit. Plundered war booty was often sold at auction. The most notable auction in Rome was held in 193 A.D. when the Praetorian Guard put the whole empire up for auction. After killing the previous emperor, the guards announced they would appoint the highest bidder as the next emperor. Didius Julianus outbid his competitors, but after two months he was beheaded by Septimius Severus who seized power. (A winner's curse?) In China, auctions were used as early as the 7th century A.D. to sell the belongings of deceased Buddhist monks. In colonial America auctions were used to liquidate inventories, unload importers' unsold items at the end of the season, and sell secondhand furniture, farm equipment, and animals. Evidently the auction was considered a disreputable way of selling goods since the owner's name was usually concealed. The most infamous auctions in American history were the slave auctions held before the Civil War.

[a]Unless otherwise noted, the historical facts presented are from Ralph Cassady, Jr., *Auctions and Auctioneering* (Berkeley: University of California Press, 1967), Chapter 3.

[b]Martin Shubik, "Auctions, Bidding, and Markets: An Historical Sketch," in R. Engelbrecht-Wiggans, M. Shubik, and R. Stark, eds., *Auctions, Bidding, and Contracting: Uses and Theory* (New York: NYU Press, 1983) p. 39.

[c]Paul Milgrom, "Auction Theory," p. 1.

their inputs with their suppliers and most people negotiate the price of a car or a house. While negotiated prices allow all aspects of the product and situation to be taken into account, they can be expensive and time-consuming, as different offers and counteroffers must be considered one at a time.

The auction mechanism falls somewhere in between posted and negotiated prices. In auctions, sellers set the rules and prices are determined by competition among potential buyers. Auctions are more flexible than posted prices. Since the price in an auction is set at the same time the object is sold, it reflects current demand conditions, the latest information, and the tastes of the particular consumers who are bidding. This flexibility is important because a common feature of the diverse items sold at auction is their uniqueness. No two oil paintings are the same even if painted by the same artist; Treasury bills sold today differ from those sold yesterday because of constantly changing information about fiscal and monetary policy as well as other economic factors. Auctions allow prices to reflect the unique aspects of goods being sold.

Auctions also differ from negotiated prices—they are less time-consuming than negotiations because the seller can compare the offers of competing buyers simultaneously rather than having to consider each offer one at a time. More importantly, once the rules of the auction are agreed on, sellers remain passive while the buyers determine the price; in other words, sellers cannot haggle with buyers as they can in negotiations. The seller's preferences only come into play when the rules of the auction are set. These auction rules serve as a commitment on the part of the seller to behave in a certain way; they also restrict the kind of offers buyers can make. So in auctions, both buyers and sellers are more constrained than in negotiations.

The rules of an auction also show exactly how a price will be determined so that demand equals supply. For example, the rules may say the winner must pay an amount equal to the highest bid. In other markets, where price adjustment is not so clear, economists have found it helpful to think about the adjustment in terms of auctions. For example, in a textbook competitive market, a hypothetical "Walrasian auctioneer" is thought to call out a price for each good, and then market participants tell him how much they demand and how much they want to supply of each good at that price. The "auctioneer" then adjusts prices—up for goods whose demand exceeds supply and down for goods whose supply exceeds demand—and the whole process continues until the market reaches equilibrium where supply equals demand, at which time trade occurs. No one believes such an auctioneer exists, but the apparatus gives economists a way of visualizing how prices move to their equilibrium levels.

Types of Auctions. Auctions are not all alike. Actually the word itself is something of a misnomer. *Auctio* means increase, but not all auctions involve calling out higher and higher bids. Auctions may take one of two basic forms, oral or sealed bid. In oral auctions, bidders hear one another's bids as they are made and can make counteroffers; each bidder knows how many others are bidding. In sealed bid auctions, bidders simultaneously submit one or more bids to the seller without revealing their bids to one another. In this case, the bidders do not necessarily know how many other bidders there are.

Two common types of oral auctions are the English and Dutch auctions. The *English* auction is the most common and well-known. The auctioneer raises the price until only one bidder remains—he wins the good at the price he has bid. In the *Dutch* auction (used to sell tulip bulbs in Holland and fish in Israel), the auctioneer calls out a high price and then continuously lowers the price until some bidder stops him and claims the good at that price.[3]

[3]Another type of oral auction, the double auction, is used on the New York Stock Exchange where many units of a good are auctioned at one time. Bids and offers are called out freely and anonymously and can be accepted immediately so that the market is continuously clearing.

In each of the different sealed bid auctions, the highest bidder wins, but the amount he has to pay differs. Consider a sealed bid auction of a plot of land and suppose three bids are received: $100,000, $99,000 and $98,000. In the *first price* auction, what you bid is what you pay if you win, so the bidder who submitted the $100,000 bid wins the land and pays $100,000. In the *second price* auction, again the highest bidder wins but he pays only the amount of the highest rejected bid, which is $99,000. Why might a seller use a second price auction? From this example, it appears he would always get a higher price by using the first price auction, but this is not true. The bidders know what kind of auction they are involved in and adjust their bids accordingly. Bidders tend to place higher bids in the second price auction than in the first price auction, so on the face of it, it is not clear which auction gives the seller the highest revenue. Still, the second price auction is rare.

Sealed bid auctions are also used to sell several units of a good at one time, such as tracts of land or Treasury bills. Though they appear more complex because multiple units and bids are involved, they are basically generalizations of the single unit first price and second price auctions. To see how these auctions work, suppose a seller auctions three identical plots of land, and he gets the following bids:

A submits 3 bids—$100,000 for one plot, $95,000 for an additional plot, and $92,000 for a third plot

B submits 3 bids—$99,000 for one plot, $98,500 for an additional plot, and $95,500 for a third plot

C submits 1 bid—$98,000 for one plot

The *discriminatory* sealed bid auction is like a first price auction: the highest bidders win and the winners pay what they bid. Each unit could be sold at a different price. In our example, the highest bids (and therefore, the prices paid) and the winners are: $100,000 from bidder A, $99,000 from bidder B, and $98,500 from bidder B again. Clearly, if two bidders win the same number of units they need not pay the same total amount for their winnings.[4]

In the *uniform* auction the units are all sold at the same price which is equal to the highest rejected bid, as in the second price auction. Since the winning three bids are $100,000, $99,000, and $98,500, the highest rejected bid is $98,000, so bidder A wins one plot of land and bidder B wins two plots, and they both pay $98,000 per plot.

During the 1960s a hot debate developed about whether U.S. Treasury bills should be sold in a uniform auction or a discriminatory auction. (See HOW TREASURY BILLS ARE AUCTIONED TODAY.) Proponents of the uniform price auction claimed that, because it was a simpler auction, bid preparation would be less costly, so more bidders would participate. This would lead to a larger volume of bills being sold, more efficient allocations of the bills, and higher bids. Proponents of the discriminatory auction claimed the government would obtain higher revenue via price discrimination.

If the type of auction had no effect on the way people bid, then the discriminatory auction would always yield the greater revenue. But buyers do bid differently depending on the rules of the auction. Economists are developing theoretical models of auction markets to examine the way bidders behave under various auction rules. Their results shed some light on why certain auctions are more common than others, what auction rules will generate the most revenue

[4]The Federal Reserve uses the discriminatory auction when it engages in short-term repurchase agreements with dealers. The auction sets the interest rate the Fed will earn on securities it purchases to temporarily increase the banking system's reserves. The Fed's outright purchases and sales for the System Account also usually occur through an auction with security dealers. See *The Federal Reserve System Purposes and Functions* (Washington D.C.: Board of Governors of the Federal Reserve System, 1984) pp. 38-43.

How Treasury Bills Are Auctioned Today

Each week the U.S. Treasury uses the discriminatory auction to sell Treasury bills to major buyers. On Tuesday the Treasury announces, via the Federal Reserve Banks, the amount of 91-day and 182-day bills it wishes to sell on the following Monday and invites tenders (bids) for specified amounts of these bills. Tenders are due by 1:00 p.m. Eastern time on the Monday after the announcement, and the Treasury usually publicizes the results later that afternoon. The bills are issued to the successful bidders on Thursday.

Two different types of bids can be submitted in the T-bill auction: competitive and noncompetitive. Competitive bidders include money market banks, dealers, and other institutional investors who buy large quantities of T-bills. The tenders they submit indicate the amount of bills they wish to purchase and the price they are willing to pay. They are permitted to submit more than one tender. Noncompetitive bidders are usually small or inexperienced bidders who indicate the amount of bills they want to purchase (up to $1,000,000) and agree to pay the quantity weighted average of the accepted competitive bids.

After all bids are in, first the Treasury sets aside the amount of bills requested by the noncompetitive bidders. The remainder is allocated among the competitive bidders, beginning with those who bid the highest price, until the total amount is issued. The price paid by the noncompetitive bidders can then be calculated based on the competitive bids that were accepted.[a]

The Treasury bill auction is more complicated than the standard discriminatory auction since the noncompetitive bids are satisfied in full. Consequently, when submitting their bids, the major buyers do not know the exact amount being auctioned to them. During 1987, an average of around $14 billion of Treasury bills were auctioned each week.

[a]See James F. Tucker, *Buying Treasury Securities at Federal Reserve Banks* (Federal Reserve Bank of Richmond, February 1985) for further details.

under different circumstances, and the crucial role of information.

PLAYING THE GAME: HOW ECONOMISTS MODEL AUCTIONS

In general, economists model the auction as a game with the bidders playing against each other. The point of the game is to win the object at the lowest possible price; each bidder devises a strategy with this in mind. The bidder's choice of strategy depends on what information the bidder has. Some information is available to all the players, like the rules of the particular auction being held. But each bidder also has *private information* about how she values the object—that is, information that only she knows. It is precisely because the bidders have some private information that sellers use an auction to set the price in the first place. If the seller knew each bidder's valuation he could just set the price of the object being sold at the highest valuation and not bother to hold the auction. The role of private information is crucial to understanding how auctions work. The assumptions made in the theoretical models about the nature of this private information range along a broad spectrum.

Independent Private Values. At one end of the spectrum, models assume that each bidder knows for certain how she values the object and that this information is totally private. The bidder's valuation of the object reflects her individual tastes; only she knows what that value is, and each bidder can have a different value. Suppose a painting is being auctioned to bidders who just want it because it is beautiful and not because they plan to sell it. (A museum might be this kind of bidder.) Then each bidder knows for certain what the painting is worth to herself but not to the other bidders, and what other bidders know about the painting will not affect her own valuation—these bidders are said to have *private values*.

Even though the bidders have private values,

each would like to know how the other bidders value the item (that is, their private information), because this would reveal something about how they are likely to bid. When the bidders' values are *independent*, then the value one bidder places on the painting is not systematically related to the values the others place on the object. In this case, a bidder's own valuation of the painting tells her nothing about the other bidders' valuations and so nothing about how they will bid.

Common Values. Models at the other end of the spectrum assume that the object being auctioned is worth the same to all bidders, but they are unsure of this value. Bidders have private information that tells them something about this true market value of the object, although not enough to be certain. For example, when the government announces a lease sale of oil and gas deposits on offshore public lands, it lets firms use seismic surveys and off-site drilling to gather information about the tracts. So different potential buyers may have different information about the market value of the tracts when it comes time to bid. The right to extract the deposits is worth the same thing to each bidder—the market value of the oil or gas actually in the land—so the bidders have common values. At the time of bidding no bidder knows this value for sure and each makes an estimate of the value based on his private information.

As in the private values model, a bidder in the common values model would like to find out what private information the other bidders have, because it would tell him something about how they are likely to bid. But, unlike the private values model, finding out their private information would also reveal something more about the likely market value of the object, which is precisely the value he is trying to estimate. Learning about another bidder's estimate, which reflects that bidder's private information, will affect a bidder's own estimate of the object's market value. Unlike the private values model, a bidder's beliefs about the value of the object can change during the course of an auction as he sees how other bidders are bidding.

Because the bidders in the common values model are unsure about the true value of the object, they are subject to the *"winner's curse."* Suppose one bidder estimates an antique chair is worth $500 based on his private information, but all the other bidders estimate its value at no more than $400. If that one bidder offers $500, he will win the chair. But by bidding his estimate, the winner is cursed! Winning conveys the message that every other bidder made a lower estimate of the chair's value, and so, on average, the winner who has bid his estimate will pay more than the chair is worth on the open market.

Bidders can avoid the winner's curse by bidding less than they think the object is worth. When there are fewer bidders, a bidder can shade down his bid more without affecting his probability of winning, because there is less chance that someone else's bid is just below his. So the seller can expect a lower price when there are fewer bidders.

Models and Reality. The independent private values model and the common values model describe extreme situations. Most real life situations are not so simple. For example, in an art auction, many bidders care about the painting's resale value as well as its personal value. Therefore, their values are neither private nor independent. Likewise, in the mineral rights auction, the value of the minerals is related to how efficiently the firm extracts them—the amount of recoverable minerals may differ for each firm bidding and so the value of the extraction rights to each firm is no longer a common value.[5]

[5]Although most theoretical research has concentrated on the polar case models, Paul Milgrom and Robert Weber have analyzed a model that includes the independent private values model and the common values model as special cases. See Paul Milgrom and Robert Weber, "A Theory of Auctions...."

A SELLER NEEDS TO KNOW THE BIDDERS' STRATEGIES...

When a seller gets ready to put his antique car or his plot of land up for auction, he has to decide which kind of auction can be expected to give him the highest price, and this will depend on how bidders behave. Analyses of theoretical models of auctions show that several factors will affect the expected price, such as whether bidders are more likely to have private values or common values, and how willing the bidders are to risk not winning the object.

Dutch and First Price Auctions. In theory, bidders behave the same way in the Dutch auction, where prices are called out in descending order, as they do in the sealed bid first price auction. So it does not matter which of these two types the seller chooses, regardless of whether the situation is an independent private values one or a common values one.[6]

A bidder follows the same strategy in the Dutch and the first price auctions because in both auctions he makes the same decision based on the same information. He knows that if he wins he has to pay what he bid, and that he wins only if he bids higher than everyone else. But he has to decide what to bid without knowing what the others are going to do. It might seem that the auctions should differ, since a bidder learns something about the other bidders' valuations during the course of the Dutch auction but not during the first price auction. But the kind of information he learns had already been incorporated into the strategy he chose at the beginning of the Dutch auction, and it is the same information he uses when choosing his strategy in the first price auction. In the Dutch auction, a bidder selects a cutoff price at which he will claim the object so long as no one else has already claimed it. As the auctioneer lowers the price, the bidder hears prices he knows are higher than other bidders' cutoff prices. But this information does not lead him to change his own cutoff price because he chose it understanding that he wins only if the other bidders have a lower cutoff value. Likewise, in the sealed bid first price auction, he selects a price knowing that it will win the object only if others have selected a lower price. Therefore, in both the Dutch auction and the first price auction, all the bidders will have the same strategy. They shade down their bids slightly below their valuations since in these auctions winners pay what they bid.

English and Second Price Auctions. The choice between the English and second price auction, on the other hand, does depend on whether the bidders know their own private values or bidders are unsure about the single common value of the item. In a situation when bidders have independent private values, both auctions yield the same outcome. In the English auction, the bidder keeps raising his bid until the price equals the value of the object to him, or until he is the last remaining bidder. Once the price equals the second highest valuation, the bidder with the second highest valuation stops bidding. The remaining bidder (who has the highest valuation) can claim the object by bidding only very slightly more than the second highest valuation. In the second price auction, the bidder simply submits a bid equal to what the object is worth to himself, since if he wins, what he pays is beyond his control anyway. Therefore, in either auction when there are independent private values, the winner is the bidder with the highest valuation and the price he pays is equal to the second highest valuation.

In a common values situation, where bidders are unsure of the value of the object being auctioned, the English and second price auctions no longer lead to the same outcome. This is because in the English auction a bidder gains two types of useful information by observing the bids of others (information he would not know at the start of the auction). He sees how many bidders

[6]However, Paul Milgrom, "The Economics...," p. 274, reports that Cox, Roberson, and Smith (in press) have experimental evidence that seems to refute this.

have fallen out of the auction (since they have lower valuations than the price being called) and he sees at what prices these bidders have fallen out. If a bidder had a reasonably high estimate of the value to start with, he gains confidence in this estimate as the bids go up, especially if many people are still in the bidding; this weakens the winner's curse and allows more aggressive bidding than in the sealed bid auction. So the price paid is likely to be higher in an English auction than in a sealed bid second price auction.

...SO HE CAN CHOOSE THE AUCTION WITH THE HIGHEST EXPECTED REVENUE.

The seller now knows that regardless of the type of information bidders have (independent private values or common values), Dutch auctions and first price auctions are expected to generate the same revenue. He also knows that with independent private values, English and second price auctions yield the same expected revenue, but with common values, the English auction is expected to be better. (See EXPECTED REVENUE DEPENDS ON TYPE OF VALUES...)

What about the Dutch and first price auction as compared with the English and second price auction? One important factor in determining which auction yields the highest revenue is how the bidders feel about the risk of losing. While each bidder in an auction would like to win, *risk averse* bidders tend to up their bids so they will be more likely to win, while *risk neutral* bidders do not.

Risk and Independent Values. In fact, in the independent private values model when all bidders are risk neutral, the Dutch and first price auctions give the seller the same revenue, on average, as the English and second price auctions. That is, while the prices are not always exactly the same, they are the same on average over a series of auctions. (This result is not obvious and remained obscure long after being proved.[7])

If, on the other hand, bidders are risk averse, then the first price auction (and therefore the Dutch) gives greater expected revenue than the second price and English auctions. In either the second price auction or English auction, risk averse bidders find it best to bid the same way they would if they were risk neutral. But in the first price or Dutch auction, risk averse bidders find it better to bid higher than they would have if they were risk neutral, as a kind of insurance against losing. (They still bid less than their valuations.) Since with risk neutral bidders, the expected revenue was the same in the four auctions, it follows that with risk averse bidders, the

[7]For our more technical readers: Bidders can be thought of as choosing, through their actions, a probability of winning and a corresponding expected payment. The revenue equivalence result hinges on the fact that, in equilibrium, the probability a bidder with a given valuation wins, is the same across all auctions in which the winning bidder has the highest valuation. In the independent private values model, all four auctions—Dutch, English, first price, and second price—have this trait. See Paul Milgrom and Robert Weber, "A Theory of Auctions...," pp. 1092-1093.

Expected Revenue Depends on Type of Values...

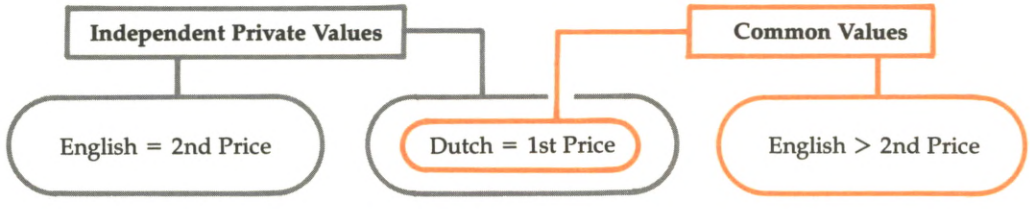

first price and Dutch auctions yield greater expected revenue than the second price and English auctions.

Risk and Common Values. Now suppose once again that bidders are risk neutral but that common values describe the situation.[8] In this case, the English auction yields the highest expected price and revenue, then the second price auction, and finally the Dutch and first price auctions. This may explain the popularity of the English auction.

We can rank these auctions using the fact that the more the price paid by the winning bidder is linked to the value estimates of the other bidders, the higher this price is expected to be. The expected price in the English auction is dependent on all the non-winners' value estimates, since the winner observes the prices at which all the other bidders have dropped out, and bases her winning bid on this information. In the second price auction, the winning price is linked to only one other value estimate—the second highest estimate—since the winning bidder must pay a price equal to the second highest submitted bid. So the expected revenue in a second price auction is less than that in an English auction. In the first price and Dutch auctions, the winning bid is not linked to any other bidder's value estimate, and these auctions yield the lowest expected revenue.

Unfortunately, if bidders are risk averse, we can no longer predict which auction yields the highest expected revenue in situations with common values or dependent values. While we know that the expected revenue from the Dutch and first price auctions is the same, and the expected revenue from the English auction surpasses that of the second price auction, the complete ranking depends on the degree of risk aversion of the bidders and on how correlated their valuations are. (See ...AND ON BIDDERS' RISK CHARACTERISTICS.)

Implications for Treasury Bill Auctions. These revenue results shed some new light on whether the U.S. should continue to use a discriminatory auction to sell Treasury bills. If the T-bill auction could be described by an independent private values model, then if bidders are risk neutral it does not matter which type of auction is used, and if bidders are risk averse the government earns higher revenue, on average, using the discriminatory auction (which is like a first price auction). But the T-bill auction seems to be more of a common values situation since

[8]Actually, all that is needed is that the bidders' valuations be dependent. That is, if a bidder places a high value on the object, she knows the other bidders are likely to place a high value on it too. Common values is the extreme case since all bidders are trying to estimate a single common value. (The estimates conditional on the true value may be independent, however.)

...And on Bidders' Risk Characteristics

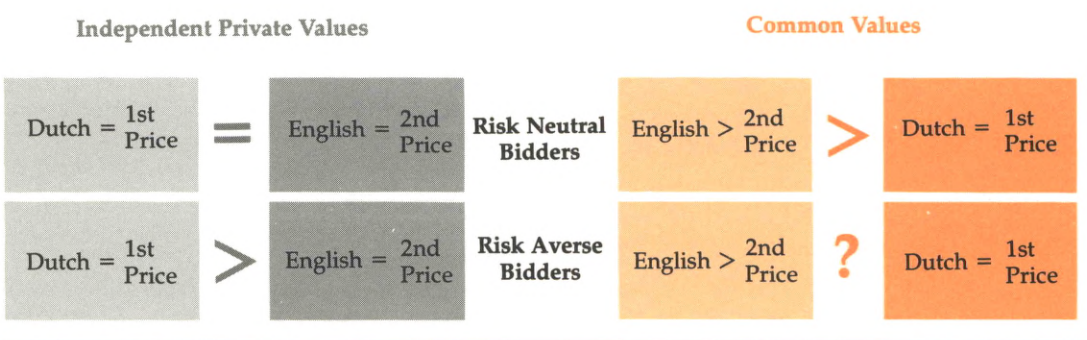

buyers are interested in the market value of the bills; so if bidders are risk neutral, the government would be better off switching to the uniform auction (which is like a second price auction). If, however, bidders are risk averse, we cannot say which auction would yield the higher expected revenue. Empirical work has compared the revenue generated by the uniform price auction which was used to sell long-term bonds in the 1970s and the discriminatory auction which is used to sell short-term and medium-term bonds.[9] The results indicate that the uniform auctions tended to yield higher revenue. This is consistent with theory when bidders have common (or dependent) values and are risk neutral.

In sum, what these results tell us is that before a seller can decide on which auction to use, he needs to find out in which situation he is likely to be. Are bidders risk averse or risk neutral? Does the situation look more like an independent private values one or a common values one? Which type of situation it is will depend on the kind of information the bidders have.

THERE'S STILL A LOT TO LEARN: THE GAVEL HASN'T FALLEN YET

Economists are just beginning to address some of the interesting questions about auctions. For example, although we know (at least in the simple models) the buyers' best strategies in an auction, we know less about when the seller will choose an auction as opposed to some other method to set the price. One reason the auction is used in a wide variety of situations is that it is efficient—the winner values the object more than any of the other bidders (and more than the seller), and he pays more than others would have paid. This means that after someone wins the object in an auction he will not be able to sell the object at a profit to someone else who participated in the auction. And the person who held the auction will not be told by a bidder afterward that she would have been willing to pay more than the winner did. So from the seller's point of view the auction can give him the highest price he can expect to receive for the object.[10]

Another area of active research concerns the seller's policy of revealing information about the object he is auctioning. Whether the seller benefits from such a policy depends upon the nature of his information. In some cases, the seller would gain, on average, from telling all the bidders his information since this would decrease the advantage some of the bidders have over others. By revealing his information the seller can weaken the winner's curse, allowing bidders to bid more aggressively. This may explain why auction houses often reveal appraisals of the objects they sell.[11] But recent work shows that adopting the policy of telling *all* kinds of information is not always in the best interest of the seller.[12] And, in fact, some government agencies conceal information about the number of firms they have invited to submit bids in their auctions.[13]

In general, most of the models studied so far have been simplifications of real life situations. Even these simple models have been difficult to analyze. But the models are becoming more realistic. Auctions with multiple buyers and sellers, like the double auction, are being

[9]See Paul Milgrom and Robert Weber, "A Theory of Auctions...," p. 1094.

[10]Another result that explains the popularity of auctions is that a seller in a poor bargaining position compared to the buyers can do almost as well as a seller in a strong bargaining position by conducting an auction. Also, a seller in a strong bargaining position sometimes will be better off selling an object by one of the standard auctions (like sealed bid, or English) than by any other method. These results are from Paul Milgrom, "Auction Theory," in Truman F. Bewley, ed., *Advances in Economic Theory* (Cambridge: Cambridge University Press, 1987), pp. 26-27.

[11]Paul Milgrom, "The Economics...," p. 287 discusses this.

[12]See Steven Matthews, "Comparing Auctions for Risk Averse Buyers: A Buyer's Point of View," *Econometrica*, 55 (May 1987), pp. 633-646.

[13]See R. Preston McAfee and John McMillan in "Auctions and Bidding," *Journal of Economic Literature*, 25 (June 1987), p. 720.

studied, as are models that assume bidders are able to collude to keep the price down. In oil and timber rights auctions the same bidders may meet again and again and so should learn more and more about their competitors—this repetition could facilitate collusion. The research that has been done suggests that some types of auctions are more susceptible to collusion than others. In one model, it is shown that collusion is easier in the English auction than in a sealed bid auction. This may explain why industrial firms, whose pool of bidders is often the same time and time again, usually solicit sealed bids.[14]

A different avenue of research being pursued involves testing the predictions of the theoretical models. Data from actual auctions are being analyzed, as are data collected from laboratory experiments. (See Herb Taylor's article in this *Business Review*.) We can expect both empirical and theoretical advances in the study of auctions in the future.

[14]This model is developed by Paul R. Milgrom in "Auction Theory."

Philadelphia/RESEARCH

Working Papers

The Philadelphia Fed's Research Department occasionally publishes working papers based on the current research of staff economists. These papers, dealing with virtually all areas within economics and finance, are intended for the professional researcher. The 21 papers added to the Working Papers Series in 1987 are listed below.

A list of all available papers may be ordered from **WORKING PAPERS,** Department of Research, Federal Reserve Bank of Philadelphia, 10 Independence Mall, Philadelphia, Pennsylvania 19106. Copies of papers may be ordered from the same address. For overseas airmail requests only, a $2.00 per copy prepayment is required.

No. 87-1 Paul Calem, "Deposit Market Segmentation: The Case of MMDAs and Super-NOWs." (Supersedes No. 86-3, "MMDAs, Super-NOWs, and the Differentiation of Bank Deposit Products.")

No. 87-2 John J. Merrick, Jr., "Volume Determination in Stock and Stock Index Futures Markets: An Analysis of Arbitrage and Volatility Effects."

No. 87-3 Richard P. Voith, "Compensating Variation in Wages and Rents."

No. 87-4 John J. Merrick, Jr., "Price Discovery in the Stock Market."

No. 87-5/R Loretta J. Mester, "Testing for Expense Preference Behavior Using Cost Data." (Revision of No. 87-5).

No. 87-6 Paul Calem and Janice Moulton, "Competitive Effects of Interstate Bank Mergers and Acquisitions."

No. 87-7 Joel Houston, "The Policy Implications of the Underground Economy."

No. 87-8 Mitchell Berlin and Paul Calem, "Financing, Commitment, and Entry Deterrence."

No. 87-9 Joel F. Houston, "Estimating the Size and Implications of the Underground Economy."

No. 87-10 Joel F. Houston, "Participation in the Underground Economy: A Theoretical Analysis."

No. 87-11 John J. Merrick, Jr., "Hedging with Mispriced Futures."

No. 87-12 Linda Allen, Stavros Peristiani and Anthony Saunders, "Bank Size, Collateral and Net Purchase Behavior in the Federal Funds Market: Empirical Evidence."

No. 87-13 John J. Merrick, Jr., "Portfolio Insurance With Stock Index Futures."

No. 87-14 John F. Boschen and Leonard O. Mills, "Tests of the Relation Between Money and Output in the Real Business Cycle Model."

No. 87-15 Anthony Saunders and Stanley Sienkiewicz, "The Hedging Performance of ECU Futures Contracts."

No. 87-16 John J. Merrick, Jr., "Early Unwindings and Rollovers of Stock Index Futures Arbitrage Programs: Analysis and Implications for Predicting Expiration Day Effects."

No. 87-17 Paul Calem, "On Estimating Technical Progress and Returns to Scale."

No. 87-18 Richard Voith, "Commuter Rail Ridership: The Long and Short Haul."

No. 87-19 Mark J. Flannery and Aris A. Protopapadakis, "From T-Bills to Common Stocks: Investigating the Generality of Intra-Week Return Seasonality."

No. 87-20 Behzad T. Diba and Herschel I. Grossman, "Rational Bubbles in Stock Prices?"

No. 87-21 Brian J. Cody, "Exchange Controls and the Foreign Exchange Market: A Model of Political Risk."

Experimental Economics: Putting Markets under the Microscope

*Herb Taylor**

There is a famous joke about a physicist, an engineer, and an economist shipwrecked on a desert island with only a can of beans to eat. "I can help open it," volunteers the physicist. "I'll start a small fire, put the can in it, and compute how long it will take for the can to explode." "Great!" says the engineer, "I can calculate the trajectory that the beans will take and where we should stand to catch them." "Wait a minute!" the economist interrupts. "You fellows are approaching this whole thing the wrong way...First, assume we have a can opener..."

Economists are notorious for making assumptions—assumptions that are at once crucial to their analysis and completely unrealistic. An economist discussing the bond market assumes that all market participants are "perfectly rational"; an economist analyzing the oil industry assumes that energy markets are "in equilibrium." Even

*Herb Taylor is a Research Officer and Economist in the Research Department of the Federal Bank of Philadelphia. The author thanks Vernon L. Smith for his informative discussion and very useful comments on an earlier draft.

economists recognize that not everyone is perfectly rational and that markets are probably never in equilibrium. Yet they stand by analyses based on such assumptions. Why?

Economists contend that it is pointless to argue over the realism of their assumptions. First of all, developing a theory always requires making some simplifying assumptions, and economists theorizing about complex human interactions are bound to make simplifying assumptions that seem exceptionally "unrealistic." But economists have a logistical problem as well. More realistic assumptions—"some people are rational" or "markets eventually settle at an equilibrium under the right circumstances"—will not improve economists' analysis unless economists can be more specific. How many people is "some"? How do the "irrational" people behave? How long is "eventually"? What are "the right circumstances" for an equilibrium? And for economists to go out into the marketplace and collect the data they need to answer these questions is a hopeless task. They cannot assemble enough information about how market participants think, choose, act, and react. They cannot control for the many factors that make one market different from another, and each market different from one day to the next. So it would seem that economists have little choice but to stick with admittedly unrealistic assumptions and hope that they are reasonable enough to produce some realistic conclusions and predictions about the way the economy performs.

Now some economists are trying a fresh approach to evaluating the assumptions economists so routinely make about the way markets operate. By constructing and observing relatively simple "experimental" markets operating under controlled conditions, they can see and test and measure the impact that different economic environments and different institutional arrangements have on market performance. Experimental economics is still young, but it has already demonstrated that taking a closer look at simpler structures opens up new ways to improve and refine economists' analyses and predictions.

ECONOMISTS ENVISION HOW MARKETS WORK IN PRINCIPLE...

Think of any good or service—Ford Escorts or visits to the dentist. Over any time period, we can observe some quantity of this product being bought at some average price. For instance, we may find that in February, 2,127 Escorts were sold at an average price of $11,359. In the mind of an economist, both the price and the quantity that we see result from the workings of "the market" for Escorts in February. But what is this "market" and how does it work to determine the price and quantity sold?

The market for a product is comprised of those considering buying the product (the demanders), those willing to provide it (the suppliers), and the social arrangements and institutions that bring them together (the market mechanism). Typically, economists' analysis of the market includes a discussion of the factors affecting the overall demand for the product, the overall supply of it, and how the two are reconciled.

Generally the public's demand for a product is held to depend on its price, prices of related products, the income level of potential customers, and their tastes and preferences. The demand curve (Figure 1) illustrates a basic idea about market demand—the higher the price of a product, the smaller the quantity consumers will want to buy.

Suppliers' willingness to make a product available is usually held to depend on the price they can get for the product, the cost of the labor, raw materials, and other factors needed to produce it, and the available technology. The supply curve (Figure 1) illustrates a basic notion about market supply—the higher the price of a product, the larger the quantity producers will be willing to make available.

Having laid out demand and supply conditions in the market, economists add the assumption that the market settles at a price which clears it of any unmet demands or unwanted supplies. This assumption gives economists a theory about the price and quantity of the pro-

duct that we observe. The price we observe is the market-clearing price; the quantity we observe is the quantity that people want to buy and sell at that price (P^* and Q^* in Figure 2).

Behind the assumption that the market for a product always clears is the economists' vision of a market as a place swarming with potential buyers and sellers, each well-informed, each operating independently, and each bidding against all others in an effort to make trades. In such a competitive environment, a product's price is persistently pushed toward its market-clearing level. A price above the market-clearing level (say at P^a in Figure 3) induces suppliers to produce more of the product than consumers want to buy, and the competition for customers forces suppliers to cut their prices. A price below the market-clearing level (say at P^b in Figure 3) makes consumers want to buy more of the product than producers are willing to make available, and competition for the relatively scarce product induces some potential customers to offer a higher price for it. Only at the market-clearing price, where consumers want to buy exactly as much of the product as suppliers want to produce, do the incentives for buyers and sellers to adjust prices disappear.

...BUT IT'S HARD TO SEE HOW MARKETS WORK IN REALITY

If market competition were really keen enough to bring markets quickly into equilibrium at the competitive price, then assuming that markets were always at their competitive equilibrium would be no problem. But real world markets are not the all-out bidding wars among teeming numbers of competitors that economic theory posits. A few products are indeed offered by large numbers of small producers, but many—like autos and breakfast cereals—are supplied by a handful of large producers, and some—like computers and fast-food hamburgers—are provided by a few small producers and one giant. Trading practices and procedures often differ from the theoretical ideal, too. In the stock market traders bombard each other with bids and

Figure 1

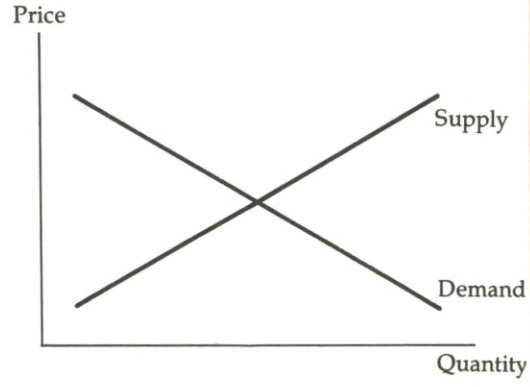

Figure 2

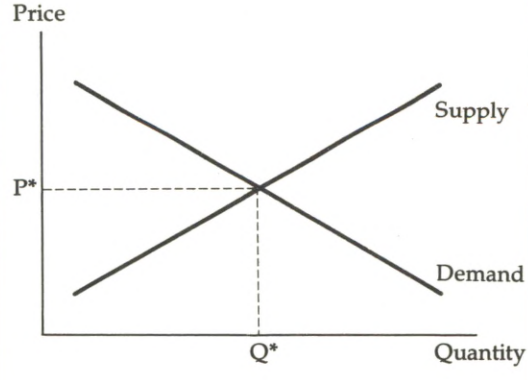

Figure 3

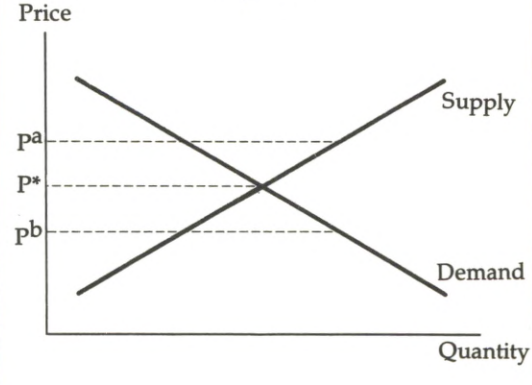

offers; but in the retail car market a salesman and a customer negotiate a price one-on-one; and in the grocery market the store clerk simply posts the price, leaving shoppers to choose between buying at that price and not buying at all. Competition may ultimately push most markets toward a competitive equilibrium. But undoubtedly some features of real world markets slow the adjustment process, and some may even block it entirely. If so, then economists could improve their analysis by developing more realistic theories about how markets work.

Unfortunately, as a practical matter, economists have not made a great deal of progress in assessing how a market's characteristics affect its performance. It is not that they haven't thought long and hard about such issues. The theory of how a monopoly producer would restrict the supply of a product to keep its price above competitive levels dates back to the 1830s. Over the years, economists have also considered how markets supplied by just a few producers (oligopolists) might behave, and they have developed many alternative theories along the way. And recently economists have gone on to develop whole new theories of how different auction formats and negotiating strategies might affect market outcomes as well.[1] But it has proven difficult for economists to assess the accuracy of these theories or to choose among the competing ones. And they have been able to offer few answers to questions like how other trading rules affect market performance, or how long it takes for a market to come to an equilibrium, or what path prices take on the way.

The problem is that economists have been trying to improve and refine their ideas about how markets work solely on the basis of what they observe in real world markets. These markets are usually so large and complicated that it is difficult—and sometimes impossible—for economists to collect the information or exercise the control that they need to test their ideas. Picture an economist trying to learn more about how product markets work by focusing on the burger business. She cannot hope to survey every fast-food producer and potential fast-food customer closely enough to get an accurate picture of supply and demand conditions in the market, so she cannot be sure what the competitive burger price would be. And she cannot add two new burger chains to the industry or temporarily switch from a posted-price to an auction format at McDonald's to see how such changes affect a market's performance.

CONTROLLED EXPERIMENTS BRING PRODUCT MARKETS INTO SHARPER FOCUS...

Economists trying to learn about how markets operate by observing existing real world markets are hobbled by their inability to observe or control the many factors at work there. Experimental economists get around the problem by setting up small markets with simple structures in which they can control for all of the relevant factors and then observe how people act in these controlled economic environments.

Setting up an Experimental Product Market. The key to researchers' control over supply and demand conditions in experimental markets is their ability to establish trade in an abstract commodity—one with no physical characteristics. Such a commodity is itself worthless; its only value is the value that the researchers induce by offering to redeem units of it for cash after the market closes. Using this "induced value" approach, a researcher can control exactly how many units of the commodity participants will want to buy and sell at any price.[2]

[1]For a more complete discussion of auctions and other trading arrangements, see Loretta Mester's "Going, Going, Gone: Setting Prices in Auction Markets" in *This Issue.*

[2]The induced value theory upon which the design of experimental markets is based is presented by Vernon L. Smith in "Experimental Economics: Induced Value Theory," *American Economic Review* (May 1976) pp. 274-279. Smith's "Microeconomic Systems as an Experimental Science," *American Economic Review* (December 1982) pp. 923-955 is a presentation of the experimental markets methodology and results which is considered the standard.

In a typical experiment, the researcher divides the market participants into demanders and suppliers. Before the market opens, she tells the demanders that any units of the commodity they buy during the trading period can be turned in for cash after the market closes. She then gives each demander a schedule indicating the redemption value of each unit he purchases in the market. Demander A's redemption schedule, for example, may indicate that the first unit of the commodity that he buys in the market can be redeemed for $.60, the second can be redeemed for $.50, the third for $.40, and so on. Likewise, the researcher informs the suppliers that after the market closes she will charge them for any units of the commodity that they sell during the trading period. She then gives each supplier a schedule indicating how much each unit he sells in the market will cost him afterwards. Supplier Z's cost schedule, for instance, may tell him that he will be charged $.20 for the first unit he sells, $.40 for the second unit, $.60 for the third, and so on. In markets for real world commodities, each market participant knows the value that he himself puts on the commodity, but not the value others put on it. To mimic this feature in experimental markets, each demander knows only his own redemption schedule and each supplier only his own cost schedule. Setting the redemption and cost schedules in this way establishes precisely the supply and demand conditions in the market. (See CONTROLLING THE MARKET WITH THE "INDUCED VALUE" APPROACH, pp. 20-21.)

Before the trading period begins the researcher also announces the trading rules. Prices may be established in one of three basic ways: auction, negotiation, or posting. Auctions allow the most interaction between buyers and sellers, negotiation somewhat less, and posted prices the least. The auction format most often used in experimental markets is the "double" auction, where both buyers and sellers are free to announce bids and offers to the market at any time. In the negotiated price format, buyers and sellers bargain with each other one-on-one. A seller may deal with a number of potential customers during the marketing period, but he must deal with them one at a time. Usually in a posted-price market, each seller decides on the price he will charge before the market opens and he cannot change it during the market period. In some posted-price experiments, the buyers of the commodity are required to decide what price they will pay in advance and they cannot change during the market period.

Once the trading rules are settled, trading begins. A trading period can last anywhere from five to twenty minutes. Usually there are eight or so participants in the market, sometimes more. Often they are college students, though working businesspeople have participated. Sometimes trading takes place in a single room; many times participants are scattered around at different locations and communicate over computer terminals. Negotiated-price markets have been conducted both using private booths to allow face-to-face contact and using telephones.

During a trading period, no money or commodities actually change hands. When a demander and supplier come to an agreement, the researcher records the price and quantity at which the transaction is completed. When the market closes, the researcher computes each participants' gains for the session. For instance, if Demander A and Supplier Z above happened to make their first transaction of the market period with each other, with Z selling A one unit at $.45, then at the end of the period A would be credited with $.15 (=$.60 − $.45) on the deal and Z would be credited with $.25 (=$.45 − $.20). If Z agreed to sell A a second unit at that price, A would be credited with an additional $.05 (=$.50 − $.45) and Z would gain with an additional $.05 (=$.45 − $.40) as well.

Once the gains from the first trading session have been computed, the researcher usually runs several more trading periods under the same market conditions to see how market behavior evolves. The researcher then may alter some aspect of the market's structure in order to observe the impact of that change on the market

Controlling the Market with the "Induced Value" Approach

Creating Demand and Supply Schedules

Participants in experimental markets are trading an abstract commodity of no intrinsic value. The researcher creates market demand for the commodity by giving each designated demander a redemption schedule and creates market supply by giving each designated supplier a cost schedule.

To create the typical looking experimental market demand schedule, the researcher could give three demanders the following redemption schedule:

Redemption Schedule

for the:	the researcher will pay you:
1st unit you buy	$.60
2nd unit you buy	$.50
3rd unit you buy	$.40
4th unit you buy	$.30

The researcher now knows that if the commodity is available in the market at a price between $.50 and $.60, each demander will make a profit on the first unit he buys, but he will lose money on the second. So, presuming that demanders prefer more money to less, each will demand exactly one unit of the commodity at a market price in that range. Market demand, then, will be exactly three units in the $.50 to $.60 price range. Similar reasoning produces the rest of the market demand schedule.

To create the typical-looking experimental market supply schedule, the researcher could give three suppliers the following cost schedule:

Cost Schedule

for the:	the researcher will charge you:
1st unit you sell	$.20
2nd unit you sell	$.40
3rd unit you sell	$.60
4th unit you sell	$.80

The researcher now knows that if the market price of the commodity is between $.20 and $.40, each supplier makes a profit on the first unit he sells, but he loses money on the second. So as long as the suppliers prefer more money to less, each will offer exactly one unit of the commodity for sale at a market price in that range. Market supply is therefore three units in the $.20 to $.40 price range. Similar reasoning produces the rest of the market supply schedule.

The demand and supply schedules created by the researcher for this hypothetical experimental market establish a competitive equilibrium price range of $.40 to $.50, indicated by the intersection of the two schedules in that range.

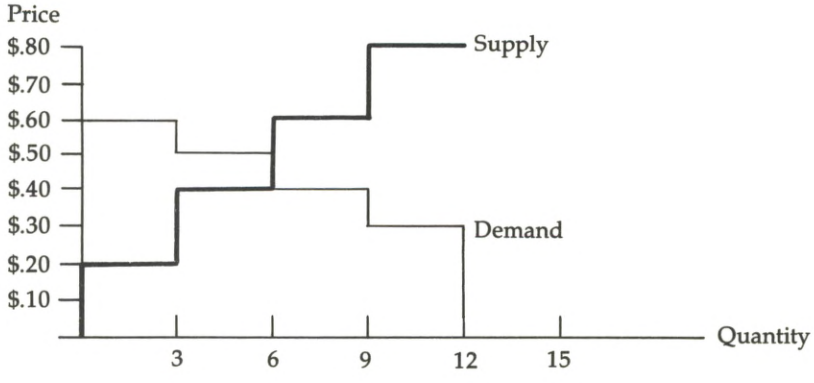

The Results of a Typical Market Experiment

In this particular experiment, the researcher, Vernon Smith, used the induced value approach to create the supply and demand conditions shown in the far left panel below. Under these conditions, the competitive market theory predicts that eight units of the commodity will be exchanged at a price of $2.10. The next two panels report what actually happened when the market was put into operation in two separate experiments of five trading periods each. In these experiments market participants bought and sold one unit of the commodity per "transaction". The number of units exchanged, as measured along the horizontal axis by the number of transactions, turned out to be somewhere between seven and nine in every period. The prices at which buyers and sellers transacted, as measured along the vertical axis, varied widely during the early trading periods of each experiment; but by the last period of each experiment, all transactions were at or near the $2.10 equilibrium price.

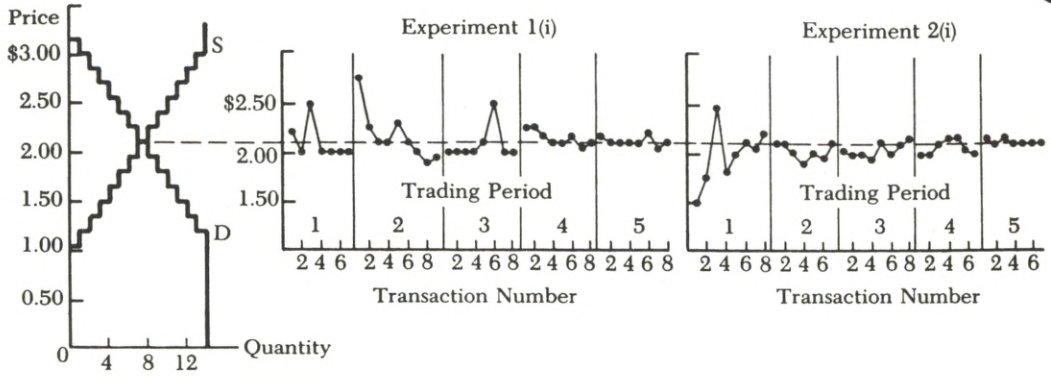

NOTE: These results were originally reported in Vernon Smith "Bidding and Auctioning Institutions: Experimental Results," *Bidding and Auctioning for Procurement and Allocation,* ed. Yakov Amihud, New York University Press (1976) pp. 43-63. The figure appears with the permission of New York University Press.

Lorem ipsum dolor sit amet, consectetur adip minimim veniami quis nostrud exercitation ull in voluptate velit esse molestaie son conseq lupatum delenit aigue duos dolor et molestai laborum et dolor fugai. Et harumd dereud facil quod maxim placeat facer possim omnis es v· atib saepe eveniet ut er repudiand sint et mol repellat. Hanc ego cum tene sentntiam, quid e amice et nebevol, olestias access potest fier a neque sonor imper ned libiding gen epular re quiet. Endium oaritat praesert cum omnin Lorem ipsum dolor sit amet, consectetur minimim veniami quis nostrud exercitati in voluptate velit esse molestaie son co lupatum delenit aigue duos dolor et mo laborum et dolor fugai. Et harumd dereu quod maxim placeat facer possim omnis atib saepe eveniet ut er repudiand sint e repellat. Hanc ego cum tene sentntiam, q amice et nebevol, olestias access potest

outcome, everything else constant. When all of the market experiments have been run, the market participants are paid their total gains from the session in cash.

Some Results of Market Experiments. Perhaps the most comforting result to come from experimental market studies is the strong tendency for auction markets to achieve the market-clearing price and quantity predicted by the competitive market model. A double auction market with a few buyers and sellers usually provides enough competition to drive the commodity price to its theoretical equilibrium price within a couple of trading periods. It seems that only a monopoly supplier can prevent the price of a product from declining to the competitive level in an auction market.[3]

On the other hand, results from experimental markets operating under different trading rules are less supportive of the competitive market paradigm. Negotiated price markets seem to converge less quickly and less directly to the competitive equilibrium than auction markets. Posted-price markets are even slower to adjust and may not converge to the competitive outcome at all. Generally when suppliers post the prices at which they will sell, the average price tends to stabilize above the competitive equilibrium level. When demanders post the prices they will pay, prices tend to stay below their competitive equilibrium level.

There is little in the way of formal theory to explain why a market's performance varies with its trading rules, but experimental economists have ventured the hypothesis that information flows play a key role. As we move from a posted-price to a negotiated-price to an auction-price format, market participants have a greater and greater opportunity to observe the terms on which others are trading and offering to trade. Access to this type of information seems to speed the market's convergence to the competitive outcome. Nonetheless, the impact of information on market performance is subtle. In one experiment, for instance, market participants' knowledge of each other's cost and redemption schedules impeded the convergence to a competitive price. In another experimental market, the researcher's release of data on suppliers' profits seemed to help them keep prices above the competitive level.

Controlled market experiments are not only helping economists isolate the conditions under which markets achieve a competitive equilibrium, they are also helping economists sort out what is happening when they do not. For instance, monopolists in experimental posted-price markets seem to achieve the higher price and restricted quantity that traditional theory suggests a profit-maximizer should, though convergence to this situation may take quite a few trading periods. Other experiments with small numbers of suppliers in non-auction markets suggest that oligopolists sometimes find some arrangement for collusion so that they can boost joint profits. For instance, when researchers ran experiments designed to mimic the major features of a barge transportation market and the market for a gasoline additive in order to address some regulatory issues, they found that adherence to certain rules for posting prices in these industries enabled suppliers to maintain higher than competitive prices.

Researchers have also used market experiments to strike out in new directions as they try to refine economists' understanding of the way markets operate. In a very practical example, one study physically separated suppliers and demanders of the commodity and then introduced a group of "middlemen" who spent one period in the suppliers' room buying and the next period in the demanders' room selling. They found that a market with a number of "middlemen" was relatively quick to achieve a

[3]This summary of experimental results is based heavily on a classic review of the literature in this area by Charles R. Plott, "Industrial Organization Theory and Experimental Economics," *Journal of Economic Literature* (December 1982) pp. 1485-1527. Another, somewhat more technical, summary is by Vernon L. Smith, "Experimental Methods in the Political Economy of Exchange," *Science* (October 1986) pp. 167-173.

competitive equilibrium. Broader in their implications are the data from experimental markets which suggest that different supply and demand conditions produce different patterns of adjustment to equilibrium: when the market demand curve has a steep slope, for instance, prices tend to start out above the equilibrium price and then decline. Economists have just begun testing some rudimentary theories that attempt to explain these kinds of patterns.[4]

...AND HELP CLARIFY HOW ASSETS ARE PRICED IN FINANCIAL MARKETS

When economists turn their attention from product markets to financial markets, considerations like buyers' tastes and sellers' operating costs move into the background and expectations play the major role. In financial markets buyers and sellers are trading IOUs—promises of future money payments—and presumably the prices at which they are willing to trade are dictated by their expectations about the value of those future payments.

According to the efficient markets theory, competition among well-informed market participants always drives a financial asset's current market price to a level which reflects the best possible forecast of its future payment stream. So the current price of a share of IBM stock, for instance, presumably would represent the best available evaluation of the dividend stream that IBM will pay in the future. Likewise, if financial markets are efficient, then the current price of an AT&T bond represents the best possible evaluation of AT&T's promise to make the interest payments and pay the face value.

At least until recently, many economists maintained that financial markets were efficient, but the October stock market crash has created some doubts. It is hard to imagine that informed market participants' best estimate of all future stock dividends could plummet by 20 percent in one day. The crash helped resuscitate a competing theory that financial markets are subject to speculative bubbles that burst. An asset's price can be bid up above its intrinsic value—the value of its expected future payout—today because some market participants believe that others will be willing to pay still more for it tomorrow. For a while this belief is self-sustaining and the market booms, but eventually participants lose faith that prices can rise further and the market crashes.

Are financial markets efficient? Do real world asset prices simply reflect a well-informed market's expectation about assets' future payout stream? Or are financial markets subject to booms and busts unrelated to changes in assets' intrinsic values? Ironically, studies of real-world financial markets cannot offer much in the way of direct answers to these important real world questions. Measuring market expectations is at the core of the problem. There are too many market participants, the possible future contingencies they must evaluate are too complex, and the constant inflow of new information changes their outlook too quickly for all of their expectations to be measured. But experimental market methods can be used to get at some answers. In an experimental asset market, the researcher can specify the payout stream of the financial asset, control the flow of relevant information to market participants, and then observe both individual and market responses. Such experiments have been run and have produced some interesting results.

Constructing an Experimental Asset Market. Experimental asset market designs are essentially multiperiod versions of experimental commodity price designs. In a typical experiment, the researcher issues each market participant some certificates which entitle the holder to dividends to be paid out at the end of each "week" of the market "year." The market "week"

[4]Vernon L. Smith, "Experimental Auction Markets and the Walrasian Hypothesis," *Journal of Political Economy* (August 1965) pp. 387-393 reports on a study of convergence paths to equilibrium in an experimental market. The study involving middlemen is reported by Charles R. Plott and Jonathan T. Uhl in "Competitive Equilibrium with Middlemen: An Empirical Study," *Southern Economic Journal* (April 1981) pp. 1063-1071.

is actually a trading session lasting several minutes; a market "year" may consist of two, three, or more market "weeks." Each participant is told what the dividend payout will be on any certificate that she holds at the end of a period, or at least told the probability distribution of the dividends—for instance, trader C may be told that for any certificate she holds at the end of "week" two she has a 50-percent chance of receiving a $1.00 dividend and a 50-percent chance of receiving no dividend. Participants are not told what payouts the other traders can expect.

The experimenter also announces the trading rules: usually experimental asset markets are organized as double auctions, just as a real-world exchange would be. Trading then begins. The experiment usually runs for several market "years" with the researcher recording all bids, offers, and transactions. The experiment can then be repeated with some alteration in experimental design in order to provide data about the impact that changing some feature of the financial environment has on the market outcome.

Experimental Evidence about Asset Market Behavior. Results from simple asset market experiments are consistent with the idea that asset markets are efficient.[5] But efficiency seems to be a fragile attribute. Studies have shown that relatively minor modifications to a simple design can easily destroy efficiency in an experimental asset market.

In the most basic asset market designs, the experimental market is run for several "years" of two or three "weeks" each, with the same weekly distribution of dividends every year. In these cases, traders tend to pick up the pattern in market prices quickly. After a few market years, each week's asset prices settle at levels consistent with the expected value of market participants' dividend streams over the rest of the market year. But in an experiment where the dividend distributions are systematically shifted from year to year, asset prices do not converge to efficient levels and fail to follow any discernible pattern. More dramatically, a batch of experiments in which the market year was simply extended to fifteen or more weeks consistently produced a speculative "boom-bust" cycle for the first couple of years. It seems that when an asset's maturity is a long way off, market participants lose sight of the dividend payments the asset is expected to yield over its lifetime and focus instead on the potential for reselling the asset at a higher price later. Only when the asset's time of maturity draws near does its expected payout become the focus of traders' attention. So there is a pronounced tendency for asset price bubbles to arise early in the market year, and for these bubbles to burst at the end as prices plunge to the efficient market price. In some of these experiments the subjects were businesspeople, not students, suggesting that it is lack of experience with a particular market situation, not an overall lack of business experience, which contributes to the speculative market behavior.[6]

Variations in experimental asset market design have produced some other interesting pieces of evidence about the way financial markets work. In one set of experiments, some traders were given "inside information" about what future

[5]Two frequently cited studies of experimental asset markets are: Robert Forsythe, Thomas R. Palfrey, and Charles R. Plott, "Asset Valuation in an Experimental Market" *Econometrica* (May 1982) pp. 537-567; and Daniel Friedman, Glenn W. Harrison, and Jon W. Salmon, "The Informational Efficiency of Experimental Asset Markets" *Journal of Political Economy* 92(3) (1984) pp. 349-408. Both articles lay out their methodology very clearly and both present results which support the efficient markets hypothesis.

[6]The nonconvergence results are reported by Arlington W. Williams and Vernon L. Smith in "Cyclical Double-Auction Markets with and without Speculators," *Journal of Business* (January 1984) pp 1-33. The boom-bust cycles are reported by Smith, Gerry L. Suchanek, and Williams, "Bubbles, Crashes and Endogenous Expectations in Experimental Asset Markets," Working Paper No. 86-2, Department of Economics, University of Arizona (forthcoming in *Econometrica*). After the recent stock market crash, Professor Smith's experimental asset market work was discussed in *The Wall Street Journal* (November 16, 1987) p. 51.

dividends on the certificates would be. The researchers found that such information was quickly reflected in the asset's market price. In several experiments, futures markets were added to allow traders to buy and sell certificates for delivery one or two "weeks" in the future. Here researchers found that the addition of futures markets reduces price volatility and speeds the convergence to efficient pricing in the spot market. In another interesting twist, market participants were prescreened and divided into two groups, more risk averse and less. Each group participated in an experimental asset market of identical design. From this experiment, the researchers concluded that less risk averse traders, those who might be termed speculators, make prices more volatile, but also help the market achieve an efficient asset price more quickly.[7]

CONCLUSION

Experimental economics—observing the behavior of subjects in controlled market environments—is giving economists the opportunity to test their assumptions and theories about market outcomes in ways that the more traditional studies of "real world" markets cannot. Much of the experimental work that has been done so far is supportive of traditional economic theory. For instance, economists' standard assumptions that product markets are competitive and that asset markets are efficient are consistent with much of the evidence from experimental markets. On the other hand, experimental work has also demonstrated that there are some important gaps and shortcomings in standard economic theory. Observing the behavior of experimental markets underscores the fact that market adjustments are not always quick, smooth, or certain.

Taking a more positive perspective, experimental economics not only points out the need to develop economic theory, but also helps economists frame new theories and provides the tools for testing them. For instance, market experiments have demonstrated that market information influences individuals' expectations and decisions in complex ways. But experiments have also provided some data to help refine economists' theories about individuals' expectations formation and decisionmaking processes. Some of this work has taken a look at some of the traditional expectations hypotheses in economics; some has tapped into psychologists' and other social scientists' theories of learning and decisionmaking.[8]

Experimental economics is a relatively new tool that researchers have developed to take a closer look at the way markets operate. So far they have put some relatively simple market structures under this new "microscope," but seeing even these simple structures up close is changing their perspective on how real world markets function.[9]

---

[7]Charles R. Plott and and Shyam Sunder, "Efficiency of Experimental Security Markets with Insider Information: An Application of Rational Expectations Models," *Journal of Political Economy* (1982) 90(4) pp. 663-698 investigate the impact of inside information. The previously cited studies by Forsythe, Palfrey and Plott and by Friedman, Harrison and Salmon introduce futures markets to their experiments. The role of risk-aversion was addressed by James S. Ang and Thomas Schwarz, "Risk Aversion and Information Structure: An Experimental Study of Price Variability in the Securities Markets," *Journal of Finance* (July 1985) pp. 825-844.

[8]Arlington W. Williams, "The Formation of Price Forecasts in Experimental Markets," *Journal of Money, Credit and Banking* (Februaŕy 1987) pp. 1-18, reports on attempts to survey market participants' expectations and model them directly. The evidence of adaptive expectations formation that Williams found is consistent with the notion that repetition of a market situation facilitates achieving the efficient markets outcome. The *Journal of Business* 59(4) pt. 2 (October 1986), in a special issue containing the proceedings of a conference entitled "The Behavioral Foundations of Economic Theory," gives some indications of how experimental economists' work ties into that of other behavioral scientists.

[9]Alvin E. Roth, "Laboratory Experimentation in Economics," *Economics and Philosophy* 1986(2) pp. 245-273 presents an enlightening perspective on the potential contributions of experimental economics as well as a thought-provoking discussion of some recent experimental results. Ken Binmore, "Experimental Economics," *European Economic Review* (1987) pp. 257-264 makes a thoughtful case for the usefulness of experimental economics.

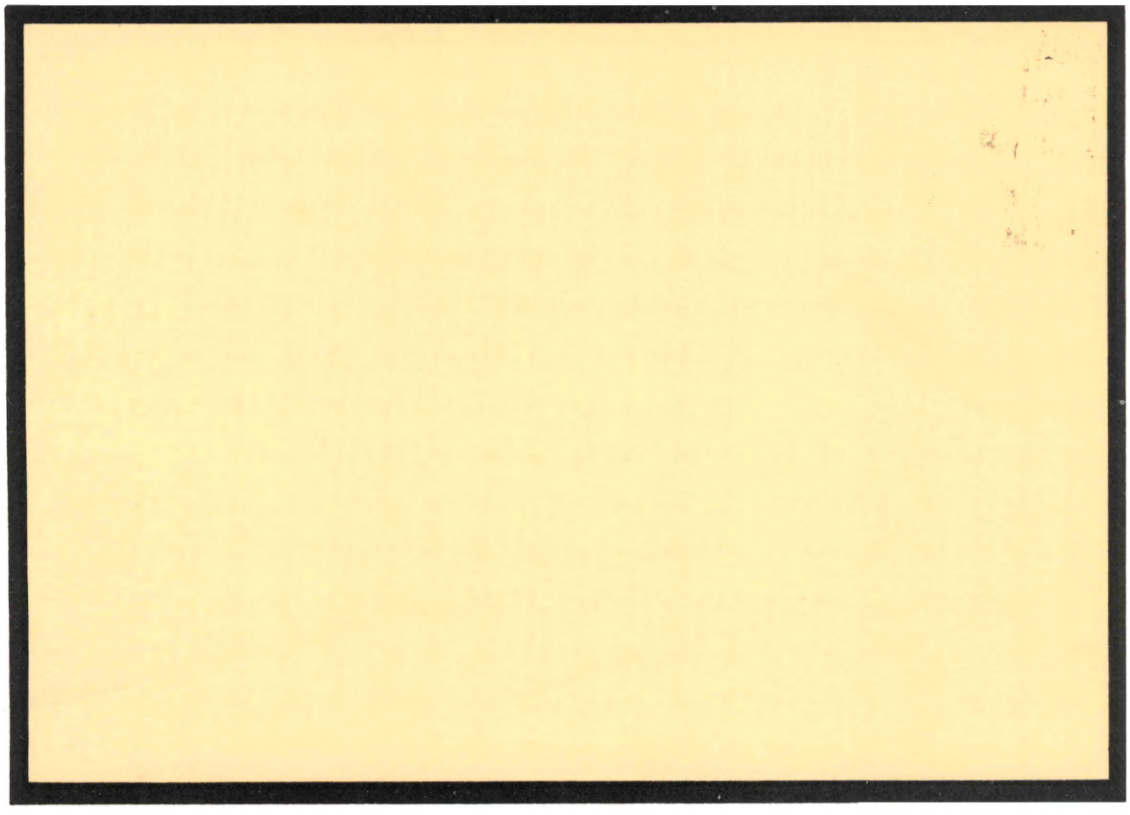

BUSINESS REVIEW Ten Independence Mall, Philadelphia, PA 19106-1574