

How Auctions Work for Wine and Art

Orley Ashenfelter

At the first wine auction I ever attended, I saw the repeal of the law of one price. The wine being sold was Chateau Palmer, a red Bordeaux, from the glorious 1961 vintage. The results of the auction I watched are contained in Table 1. Each lot (of 12 bottles) of wine for sale was identical, and three lots were sold in a matter of seconds by Sotheby's according to the usual English auction practice. The first lot fetched 920 pounds, the second lot fetched 800 pounds, and the third lot fetched 700 pounds. The first bidder paid 31 percent more for his 1961 Chateau Palmer than did the third bidder!

I was, of course, stunned to see such enormous price variability for identical objects sold at the same time and place. Even more amazing, as I looked around the auction room it seemed obvious that no one else in the room thought anything was amiss. The law of one price was repealed and no one even seemed to notice!

This empirical surprise led me to begin collecting data on wine auctions, to interview auctioneers, and even to buy a little wine.¹ In the meantime I have also had the opportunity (with John Abowd, 1988) to observe and collect data on the auction sale of impressionist and contemporary paintings. This paper reports on some of the empirical regularities that I and others have observed in the actual operation of the auction markets for these items. In view of the rich and diverse array of theoretical literature on auctions it seems high time economists began to spell out precisely what facts it is meant to explain.²

¹Collecting data is an expensive hobby; to support it I began publication of the results of the world's wine auctions in a convenient financial format. See *Liquid Assets: The International Guide to Fine Wines*, 169 Nassau St., Princeton, NJ 08542.

²Recent empirical analyses of oil leases by Hendricks and Porter (1988) and of highway bidding by Thiel (1988) show how rich is the potential for the study of actual auction outcomes.

Table 1
Typical Price Sequences for Identical Wines
Sold in a Single Auction
(in £ sterling per dozen bottles)

<i>Wine Type</i>	<i>Chateau Palmer 1961</i>		<i>Croft (Port) 1927</i>		<i>Chateau Margaux 1952</i>		<i>Quinta de Noral (Port) 1934</i>	
	<i>Lot Size</i>	<i>Price</i>	<i>Lot Size</i>	<i>Price</i>	<i>Lot Size</i>	<i>Price</i>	<i>Lot Size</i>	<i>Price</i>
Lot 1	12 bts	920	12 bts	800	12 bts	480	10 bts	400
Lot 2	12 bts	800	12 bts	800	12 bts	480	12 bts	500
Lot 3	12 bts	700	12 bts	750	12 bts	480	12 bts	500
Lot 4			12 bts	650	24 bts	480	12 bts	480
Lot 5			12 bts	650	24 bts	480	12 bts	480
Lot 6			12 bts	650	20 bts	480		
Lot 7			12 bts	650				

Source: Author's tabulation of the results of Sotheby's Auction of Finest and Rarest wines, December 11, 1985.

How "English Auctions" Really Work

Many people think they understand the rules of an English auction because they are so commonly used.³ Sotheby's, Christie's, Phillips, and the other English auction houses have invented and refined these rules over two centuries, and they are now common in many other parts of the world. It is well known that in an English auction the bidding begins low and edges upward as bidders escalate their bids. When the bidding stops, the item for sale is said to be "knocked down" or "hammered down." The price at which an item is knocked down or hammered down is called the "hammer price."

What is not so well understood is that the items knocked down have not necessarily been sold. Here is the reason. The seller will generally set a "reserve price," and if the bidding does not reach this level the item will go unsold. Auctioneers say that an unsold item has been "bought in." (This terminology is somewhat misleading since unsold items are rarely, if ever, bought by the auction house.) An item that has been bought in may be put up for sale at a later auction, sold elsewhere, or taken off the market. John Abowd and I (1988) found that about one-third of the Impressionist paintings put up for sale in the early 1980s did not find buyers. In wine auctions, on the other hand, the typical "buy in" rate ranges from 5 percent to 10 percent. It is my impression that the typical buy in rates for other auction items—European paintings, silver, furniture, and jewelry—fall between these extremes.

³What is called an English auction is, in fact, Roman. The word auction comes from the Latin "auctio," which means to ascend.

Auctioneers are very secretive about whether and at what level a reserve price may have been set, and there is a real auctioneer's art in getting the bidding started on each item without revealing the reserve price. For example, the auctioneer may have to accept and announce fictitious bids "off the wall" or "from the chandelier" to start the "real" bidding. Bids from off the wall are legally being placed on behalf of the seller. At the same time, sellers are forbidden by contract with the auctioneer from bidding in the auction. This is the protection that the auctioneer offers to the prospective buyers to ensure that they are not being artificially "bid up."

If you sit through an auction you will find that every item is hammered down and treated as if it were sold.⁴ Only after the auction does the auctioneer reveal whether and at what price an item may have actually been sold. In short, the auctioneers do not reveal the reserve price and they make it as difficult as they can for bidders to infer it.

Although the above description outlines commonly accepted practice in auctions, many economists who write about auctions describe them differently. For example, Milgrom (in this journal) states "the auctioneer begins with the lowest acceptable price—the reserve price—and proceeds to solicit successively higher bids from the customers until no one will increase the bid. Then the item is 'knocked down' (sold) to the highest bidder." As I have noted, real auctioneers do not reveal the reserve price in this way, and many "knocked down" items may be unsold. In another example, Graham and Marshall (1987) state that, "when the bidding stops, the auctioneer will generate a false or phantom higher bid if he feels that the high bidder is 'good for another bump'." However, inventing fictitious bids above the reserve price is certainly unethical and probably illegal, too. Since the auctioneer's rules are known to an entire array of personnel who often move on to become bidders or their agents, it would soon become common knowledge if an auction house systematically engaged in the generation of fictitious bids above the reserve price. No auction house that values its reputation—and the long run profits its reputation secures—would systematically engage in this practice.⁵

Secret reserve prices and high "buy in" rates have some interesting implications for the theoretical study of auctions. In the optimal auctions model of Riley and

⁴There are exceptions to this rule. In New York City the auctioneer is legally required to state whether an item has been sold at the conclusion of the bidding. The New York auctioneers did not reveal this information before this rule was promulgated, and they do not reveal this information in other locations except where required by law.

⁵It should be appreciated that an auctioneer faces a real tradeoff in deciding whether to follow the standard ethical auction practices. If the auctioneer gains a reputation for following these practices he receives the benefit that buyers will be willing to reveal their true valuation of an item, and this will, in the long run, lead to higher prices for the items sold. On the other hand, there are always short term gains to be obtained by abandoning the standard practices. I do not mean to suggest that all auctioneers follow the standard practices considered ethical in the auction trade, although I do believe that the large auction houses with substantial reputations follow them. Many sellers try to give the impression that they follow standard auction practices even when they do not, apparently because they believe that it increases their sale prices to uninformed buyers. Some of the most amusing examples of this practice are depicted by the home shopping clubs broadcast on late night cable television. This kind of cheating seems to be endemic to any kind of economic activity where reputations are valuable.

Samuelson (1981), for example, the reserve price serves to extract a slightly higher price from the bidder with the highest valuation of the item on offer.⁶ The reason is that in an English auction the seller only receives an amount equal to the second highest valuation placed on the object among the bidders. By setting a reserve price the seller takes a chance on extracting part of the valuation gap between the two bidders with the highest valuations in exchange for the risk of losing the sale altogether.

However, this optimal auctions model is probably not much help in understanding how reserve prices are set in most auctions. First, since it is a dominant strategy for each bidder in an English auction to bid up to their true valuation of the object, the optimal reserve price is identical no matter whether it is kept secret or not. This model therefore offers no explanation for why the reserve price should be secret. Second, in the optimal auctions model the reserve price is independent of the number of bidders. It follows that the probability that an item will be sold increases with the number of bidders. In fact, however, as I noted above buy in rates are very high for some types of items despite a large number of bidders. Moreover, buy in rates differ systematically across types of items in a manner that is almost certainly not related to the number of bidders in these auctions. It seems very unlikely that actual buy in rates can be explained primarily by the considerations important in the optimal auctions literature.

One likely explanation for the secrecy surrounding reserve prices is that it serves to thwart “rings.”⁷ There is always random variation in the interest and turnout of bidders; when the turnout is low, some sellers may prefer that their goods be bought in and offered for sale at a later date rather than risk a collusive ring bidding to depress the item’s price. The auctioneer may also engage in other practices that weaken rings. For example, the auctioneer typically does not reveal the identity of the purchaser, if there was one, and this creates strong incentives for the ring members to bid privately in opposition to the interests of the ring.

An explanation for the key determinants of the seller’s reserve price may be found in models of search (Mortensen, 1970), where the seller may expect to offer the item at auction more than once, or even to sell it privately to a dealer as an alternative. The highest observed price in a particular auction may be thought of as a “job offer” which will be accepted only if it exceeds the reserve price. In these models there is a “natural rate of unemployment” that may well be related to the “normal buy in rate” that characterizes auction markets. Abowd and I (1988) find that time-series movements in the auction market for Impressionist paintings do seem consistent with the dynamics of search models, so perhaps the other characteristics of these markets are consistent with them as well.

⁶The term “reserve price” is an unfortunate choice of words in this context. In reality virtually every seller has some price below which they would not agree to sell an object; the theory of optimal auctions indicates why, for strategic reasons, a seller should set a reserve price that is strictly higher than the minimum price for which they would sell the object.

⁷Webster defines a “ring” as, “An exclusive combination of persons for a selfish, and often corrupt, purpose, as to control the market.”

Competition Among Auction Houses

It is sometimes said that the auctioneers at Christie's (still owned by Englishmen) are gentlemen who try to act like businessmen, while the auctioneers at Sotheby's (now owned by Americans!) are businessmen who try to act like gentlemen. There is no doubt an element of truth to this characterization of the style of these two auction houses. The competition among auctioneers is more than a matter of style, however. The auction business is an interesting example of an industry where the cost of building a reputation may act as a significant entry barrier to new competitors. Surprisingly little attention has been paid in the literature about auctions to the role of the auctioneer.

In principle the auctioneer acts on behalf of the seller, but the auction house typically receives compensation from both the buyer and the seller for items that are sold. The *buyer's premium* is a percentage of the sale price paid to the auctioneer by the buyer. In most auction houses the buyer's premium is 10 percent of the sale price, and the amount is generally not negotiable. The *seller's commission* is a percentage of the sale price paid to the auctioneer by the seller. It varies with the type of item being sold, ranging from 15 percent for wine to as little as 5 percent for certain kinds of paintings. (Thus, if a single lot of wine is sold at an auction for a hammer price of \$100 the buyer will pay \$110 to the auctioneer, and the seller will receive \$85 from the auctioneer.)

If an item goes unsold, the auctioneer will receive neither a buyer's premium nor a seller's commission. To make sure the seller bears some of the cost of auctioning but not selling an item, auctioneers usually charge the seller a fee on unsold items. This fee is often a percentage of the reserve price set by the seller, which obviously gives the seller an incentive to keep the reserve price low. In addition, some auction houses will not allow a seller to put up an easily recognized item for resale until some time has passed. Sometimes it is claimed that when an advertised item goes unsold its future value will be affected. Such items are said to have been "burned."

Whether an item can be "burned" has some implications for the set of assumptions used in studying auctions. When bidders have independent private valuations of the items on offer, as in Riley and Samuelson (1981), it should not be possible for any item to be burned. The failure of an item to sell should not influence its future salability either at auction or privately. But in models where bidder valuations are correlated, as in Milgrom and Weber (1982), the failure of an item to sell is informative about the value it might achieve in another sale. If an object can be burned, the assumption of correlated valuations becomes more appropriate. However, it is not clear whether paintings really are burned, or whether this is a fiction invented to encourage sellers to be satisfied with lower reserve prices.

Most auction houses now collect a much larger part of their revenues from the buyer's premium than was true in the past. The most commonly given reason is that auction houses have been increasingly forced to bargain down the size of the seller's commission when they deal with large consignors. Since it is far more difficult for a

Table 2
Wine / Hammer Price Differences Across Auction Houses

<i>Auction House</i>	<i>% Difference from Christie's London Prices</i>			
	<i>Spring 1986</i>	<i>Fall 1986</i>	<i>Spring 1987</i>	<i>Fall 1987</i>
Sotheby's London	- 12%	0%	+ 5%	- 4%
Christie's Chicago	+ 14%	+ 18%	+ 15%	+ 4%
Butterfield's San Francisco	+ 27%	+ 21%	+ 10%	0%
Christie's Amsterdam	- 16%	- 8%	- 8%	- 8%
Christie's Geneva	+ 13%	+ 37%	+ 25%	+ 19%
Sotheby's Geneva	+ 27%	+ 34%	+ 27%	+ 30%

Source: "Liquid Assets," *The International Guide to Fine Wines*, Issue No. 4, Spring 1988.

large number of small buyers to bargain effectively, the buyer's premium has provided a new source of revenue. This suggests that real bargaining costs have changed in recent years, but I am unaware of why this might have happened.

The result of an interesting natural experiment in the variation of the buyer's premiums is contained in Table 2. This table contains a comparison of the hammer prices (that is, not including the buyer's premium) for identical wines sold at six auction houses at four points in time. In the spring of 1986 buyer's premiums were 10 percent in Sotheby's London, Christie's Chicago, Butterfield's San Francisco, and Christie's and Sotheby's Geneva auctions. The buyer's premium was 16 percent in Amsterdam at this time, and there was no buyer's premium in Christie's London auctions.

It is reasonable to consider Sotheby's London, Christie's London, and Christie's Amsterdam auctions to be conducted in the same market. The other auctions take place in markets that are divided by both regulatory and transportation cost barriers. Rational buyers will, of course, discount the price of a wine to the extent that it carries a buyer's premium. Thus, the price differences among Christie's London, Sotheby's London, and Christie's Amsterdam auctions in Spring 1986 are consistent with the assumption that these auctions are taking place in a single market and that buyers have caused buyer's prices to be equalized. Prices in the United States and Geneva, on the other hand, offered opportunities for arbitrage in early 1986.

In the fall of 1986, Christie's instituted a 10 percent buyer's premium in their London auctions. The Christie's auctioneer had earlier condemned Sotheby's for instituting a buyer's premium, and had promised not to follow suit, but some aspect of the commercial realities did not allow Christie's to keep this pledge. Most observers believed that having instituted a buyer's premium which would definitely generate some revenue, Sotheby's had then proceeded to negotiate with the larger trade-oriented sellers over the seller's commission and had threatened to take away some of

Christie's larger and more lucrative clients. Christie's put a halt to this behavior by implementing their own buyer's premium.

As shown in Table 2, the institution of a buyer's premium by Christie's London immediately resulted in a realignment of hammer prices, so that buyer's prices remained the same among Christie's and Sotheby's London auctions and Christie's Amsterdam auctions. In the meantime, buyer's prices have also converged between London and the United States, albeit much more slowly. Perhaps some arbitrage did take place. Geneva prices, on the other hand, remain remarkably high, suggesting the existence of some permanent barriers to trade between Switzerland and both the U.S. and England.

The Price Decline Anomaly

I believe it is common knowledge among auctioneers that, when identical lots of wine are sold in a single auction, prices are more likely to decline than to increase with later lots. This does not mean that price declines always occur, but they are far more common than would be expected by chance alone. Table 1 provides some examples of the pattern of results that you are likely to see in a wine auction.

A more formal analysis is contained in Table 3. I have compared the prices of every pair of identical wines sold in the same lot size in the years 1985 through 1987 in London, Chicago, and San Francisco. Although it is most common for the price to remain constant, prices are at least twice as likely to decline as to increase. This pattern is true for every auction house in London or the United States.

Table 3

The Distribution of Price Patterns for Identical Wines Sold in the Same Auction (August 1985–December 1987)

	Auction House			
	Christie's London	Sotheby's London	Christie's Chicago	Butterfield's San Francisco
Later Price Higher	271	143	90	20
Later Price Lower	628	430	183	41
Later Price Identical	1498	1073	226	39
Mean Ratio (2nd to 1st price)	.9943	.9875	.9884	.9663
Standard Error of the Mean	.00128	.00188	.00335	.0103
Number of Comparisons	2370	1646	499	100

Source: "Liquid Assets", *The International Guide to Fine Wines*, Issue No. 4, Spring 1988.

Table 3 also contains the means (and estimated standard errors) of the ratios of the second to the first prices. If prices followed a random walk these ratios would be unity, but all the ratios are less than unity, and the differences are statistically significant. However, since the cost of buying and then selling these lots of wine will rarely, if ever, be less than 15 percent of their value, these discrepancies will not support a profitable trading strategy. Still, if you attend a wine auction it is worth remembering that it does not pay to be too aggressive in bidding for early lots of a particular wine!

Most auctioneers are aware that later bidders on similar items are more likely to pay lower prices, but they are uncomfortable about revealing this information to uninformed bidders. Apparently most bidders, like economists, expect to see identical items sold at identical prices! When inexperienced bidders see exceptions to this rule they may think something fishy is going on. As a result, auctions are set up to disguise this regularity. For example, the auctioneer will usually offer smaller lots of the same item before larger lots. Since most bidders see nothing anomalous in quantity discounts, declining per unit prices seem more acceptable. Of course, the statistical analysis I reported above deletes cases where different lot sizes are being offered, and this selection bias has probably led me to underestimate the extent of the decline in prices.

Auctioneers have another device for limiting the extent to which bidders are likely to see price declines for identical items. When a series of lots of identical items is offered, the winning bidder on the first lot has the option of immediately taking all the subsequent lots at the same price.⁸ This rule has two effects. On one hand, for a bidder who wants some (but not all) of the items on offer, it increases the risk of waiting for the lower prices that may materialize with later lots. Risk averse buyers are thus forced to pay a real price for any attempt to exploit the typical pattern of price declines: they may lose the opportunity to buy any of the lots they want. Thus, this rule increases the seller's revenue so long as there are risk averse buyers. In addition, since the option to purchase several lots is often exercised, bidders will in this case see a uniform price for all items. (This is also another reason to believe that I have underestimated the extent of the anomalous declines in prices that would otherwise have taken place in Table 3.)

These anomalous price declines appear to be characteristic of many other auction markets. For example, Penny Burns (1985) compared the behavior of student subjects and a group of experienced Australian wool traders in a simulated auction environment identical to that used in the wine auctions described here. After a few "weeks" of simulated trading, the students learned to maximize their profits and market prices followed a random walk throughout trading "days." The wool traders operating in

⁸Since the buyer must exercise this option immediately, this rule does not establish an option value for risk neutral traders which could be used to explain the price decline anomaly. Here is the reason. Since the option expires as soon as it is purchased, it is only of value if it is exercised. If it is exercised, however, there will be no price decline. Thus, the existence of an option value is not consistent with the price declines. For those who like empirical tests of such theoretical arguments I might also add that Butterfield's San Francisco auctions have no "buyer's option"—but the price decline anomaly is even more severe there.

the same environment did not learn to maximize their profits and market prices in their simulated trading declined as in the wine auction results above. Burns asked the wool traders whether they had noticed the pattern of declining prices and reports that, “Despite its consistency over 15 consecutive auction sessions, they had not seen the price fall!” This suggests that in the market trading to which they are accustomed, wool traders do not consider a declining price pattern to be a signal of unexploited profits.

These results suggest that risk aversion or quantity constraints play a significant role in real auction markets. Indeed, assuming bidders are risk averse may simply be a convenient analytical device for dealing with the fact that many bidders at auctions are buying to fill orders and are effectively quantity constrained. In fact, the mechanism used in practice by the auctioneers that gives the first buyer the option to purchase subsequent lots at an identical price is clearly related to the optimal auction design suggested by Maskin and Riley (1984) when there are risk averse bidders. Theoretical work on auctions will almost certainly have to remove the assumption of risk neutral bidders if it is to explain the full range of interesting empirical results from real auctions.

The Information in Auction Results: An Externality

A young woman called me with an unusual request just before Christmas: “Professor Ashenfelter, my husband was killed in a car accident last month, and I am going to have to sell his wine cellar. A relative has offered to take the wine for \$3,000, but I am suspicious it may be worth more. Would you tell me what it is worth?”

Oh my, what a sad situation. Knowing that younger wine collectors have young (and not very valuable) wines I wasn’t quite sure what to say. So, after offering my own condolences I asked, “What sorts of names are on the bottles?”

“I have a list; there are Chateau Lafite, Chateau Latour, Chateau Mouton . . .”

These are, of course, the names of the finest red Bordeaux wines. Maybe, I thought, this young woman had a little better luck than I had first guessed. So I asked, “What vintages are on the bottles?”

“There are quite a few. The oldest is 1945, but there is 1961, 1959 . . .”

As it turned out, my estimate was that this young woman’s wine cellar would fetch about \$18,000 if it were put up at auction, about six times what her relative had offered her!

This story illustrates one way that regular auctions confer information benefits that are typically not captured by the profits of the auctioneer. In essence, the presence of an auction system provides a way for an uninformed seller to obtain approximately the market value for the items they own.

The value of a public auction system as protection for uninformed sellers has long been understood in Europe, but it is not widely appreciated in the United States. In both Sweden and Austria, for example, the auction houses are run as state-owned monopolies. The major auction houses in both the U.S. and England are certainly not

state-owned, but they are watched with some care by institutions like the Metropolitan Museum of Art, which disposes of its property only through public auctions.

My advice to the young woman whose wine cellar I appraised was, of course, to sell her wine at auction unless she could find a private buyer who would pay her at least (the certainty equivalent of) what I estimated she would receive for the wine if it were sold at auction. The information on prices I used would not exist without a system of public auctions. In other words, the auction system allows traders to make private transactions outside the auction system, but if everyone traded in this way there would be no auctions in the first place. I believe that the inability of auctioneers to capture a significant part of the benefits of the information they produce leads to less use of the auction system than is optimal for society.

Of course, some people oppose the use of the auction system as a matter of principle. For example, many wine merchants argue that the auctioneers pose “unfair competition” because they do not have to maintain an expensive inventory and therefore the prices realized at auction will be lower than the merchant’s. To an economist, such special interest pleading is hard to take seriously, but state legislators are not so immune. The result is that wine auctions are illegal except in Illinois and California.

The Information in Auction Results: Taking Advantage of It

In 1987, not long after public wine auctions had begun in California, I asked Dennis Foley of the San Francisco auction house of Butterfield and Butterfield how I could find the “real bargains” in his wine auctions. He said, “Buy the wines Robert Parker doesn’t like. Parker’s book (1985) has an enormous influence on U.S. auction prices. Wines he doesn’t rate or mistakenly rates make fabulous bargains.”

To U.S. wine buyers this advice may border on the heretical. Here is the reason. Most of the information available about the quality of wine appears in publications that are run by advertisers. Newspaper writers often accept direct and indirect payoffs for touting wines and many knowledgeable consumers are aware of this practice. (After a detailed expose by media reporter David Shaw in the *Los Angeles Times* in the summer of 1987, the *Times* fired its own wine writer!) The result is that there is a market for independent information. Robert Parker, lawyer turned wine writer, has captured that market and the absolute faith of his readers by refusing to accept handouts and payoffs.

What Parker and many other wine writers attempt to do is to infer the likely quality of a red Bordeaux, Burgundy, or California wine when it is 10 or 15 years old from the taste of the wine when it is a few months old. In other words, wine consumers in search of their liquid asset are much the same as oil drillers in search of their own liquid asset. In the future the true value of the asset will be revealed, but meanwhile the consumer must be satisfied with an estimate of its value. What Dennis Foley seemed to be telling me was that U.S. wine consumers were suffering from an extreme version of the winner’s curse. Consumers would bid on a wine in his auctions only up

to Parker's valuation of it, as if Parker's appraisals could not possibly contain error. It is precisely this kind of failure to make probability judgments that seems to be the cause of the winner's curse found in experiments by Kagel and Levin (1986).

The question is, of course, how to take advantage of the errors of these bidders. As it turns out, the theoretical literature on the information content of auctions suggests a possible way to do this. In a classic paper Wilson (1977) shows that the auction price is a better estimator of the true value of an item as the number of bidders gets larger. Since the volume of wine sold at auction in London is more than ten times the U.S. equivalent, and the number of potential bidders may be greater still, why not use the prices from London's auctions as benchmarks, buying wine in the U.S. auctions only when its price is systematically lower than in London?

If this form of arbitrage works, it will not make you (or anybody else) any money. Unless the wine is shipped to London, no one in the U.S. is likely to buy it for more than you paid for it. But if you buy the wine to drink, you may have purchased a great bargain.

Over the last year or so I have tested this strategy for wine auction buying several times. So far it has worked very well, although there are signs that enough market participants are learning about it so that it may not last a great deal longer. Here is a recent example that you might still run across: 1970 Chateau Brane-Cantenac, a red Bordeaux, sells for about \$330 per dozen bottles in London, while it regularly fetches about 20 percent less in both San Francisco and Chicago. I and several colleagues recently sampled this wine and found it to be excellent.⁹ Why is this wine so much cheaper in the United States? Well, here is what Parker (1985) says about it: "This is a distressingly poor wine, particularly in view of the vintage. The 1970 has quite a foul aroma and dirty barnyard scent. On the palate, the wine is beginning to fall completely apart. Lacking fruit and concentration, this is a vivid example of a very poorly and sloppily made wine." I can assure you that neither I nor my colleagues detected any barnyard scents in this wine and I do not believe that Parker would either.

Auctioneers and Accurate Information

The theoretical literature about auctions emphasizes that there are good reasons for auctioneers to provide truthful information about the items being sold (Milgrom and Weber, 1982). The basic idea is that revealing information tends to remove uncertainty and make low bidders more aggressive; this puts upward pressure on the bidding of others, which is in the interest of the auctioneer. It may seem surprising to some, but auctioneers do appear to act consistently with this prediction.

For example, auction houses typically go to considerable effort to estimate the price that an item offered for sale will fetch. Predicting the price at which a unique

⁹I am indebted to my colleagues Dwight Jaffee and Richard Quandt for their help in this empirical work and to Burton Malkiel who also supplied the wine.

Table 4

**Comparison of Actual and Estimated (by the Auctioneer Prior to the Auction)
Sale Price of Impressionist Paintings 1980–1982**

<i>Auction House</i>	<i>Number of Paintings Put up for Auction</i>	<i>Percent not Sold</i>	<i>Percent Sold at Price Below the "Low Estimate"</i>	<i>Percent Sold at Price Above the "High Estimate"</i>	<i>Percentage Difference Between Sale Price and Estimated Sale Price</i>
Christie's London	211	44%	18%	20%	3.9%
Christie's New York	433	30%	21%	24%	2.1%
Sotheby's London	918	34%	21%	24%	1.5%
Sotheby's New York	617	26%	19%	28%	3.6%

Source: Abowd and Ashenfelter (1988).

item will sell requires considerable expertise, so it is of some interest to see just how good these predictions are. The usual practice is for the auctioneer to provide a high estimate and a low estimate in an auction catalogue.


Taking the auctioneer's point estimate to be the midpoint of the high and low estimates, Table 4 contains the results of an empirical study of how well the auctioneers predicted the prices fetched by Impressionist paintings in London and New York in auctions held by Sotheby's and Christie's. Figure 1 provides a scatter plot of the actual prices against the auctioneer's estimates.

These data reveal some important results. First, the auctioneer's price estimates are very highly correlated with the actual prices fetched and they are very close to unbiased. Since the prices fetched are available only for pictures that are actually sold, unbiased estimates averaged over only those pictures who find a buyer will appear downward biased. Although the extent of this selection bias is very small, this is precisely what the data indicate for both auction houses operating in both London and New York. It is also worth observing that, whatever implicit rules the auction houses use for defining their high and low price estimates, the results of these rules lead to very similar points on the empirical frequency function of forecast errors.

In sum, auctioneers do seem to provide genuine expertise in predicting prices. In fact, Abowd and Ashenfelter (1988) find that the auctioneer's price estimates are far better predictors of the prices fetched than any hedonic price function presently available. Perhaps honesty is an auctioneer's most profitable policy rule.

Figure 1

The Relationship Between Predicted and Actual Prices for Impressionist Paintings



■ I am deeply indebted to several people in the auction trade who have been very generous in describing the institutions with which they work. Brian Cole, Director of Auctions at Christie's New York, Michael Broadbent, Director of the Wine Department at Christie's, and Bruce Kaiser, Director of the Wine Department at Butterfield and Butterfield, San Francisco have all been exceptionally helpful. I have also received helpful comments on an earlier draft of this paper from David Card, Robert Marshall, Robert Wilson, and the editors of this journal.

References

- Abowd, John, and Orley Ashenfelter**, "Art Auctions: Price Indexes and Sale Rates for Impressionist Pictures," unpublished manuscript, 1988.
- Burns, Penny**, "Experience and Decision Making: A Comparison of Students and Businessmen in a Simulated Progressive Auction," Vol. 3 of *Research in Experimental Economics*. Greenwich, Conn.: JAI Press, Inc., 1985.
- Graham, Daniel A., and Robert C. Marshall**, "Collusive Bidder Behavior at Single-Object Second-Price and English Auctions," *Journal of Political Economy*, December 1987, 95, 1217–39.
- Hendricks, Kenneth, and Robert H. Porter**, "An Empirical Study of an Auction with Asymmetric Information," *The American Economic Review*, December 1988, 78, 865–83.
- Kagel, John H., and Dan Levin**, "The Winner's Curse and Public Information in Common Value Auctions," *The American Economic Review*, December 1986, 76, 894–920.
- Maskin, Eric, and John Riley**, "Optimal Auction with Risk Averse Buyers," *Econometrica*, November 1984, 52, 1473–1518.
- Milgrom, Paul R., and Robert J. Weber**, "A Theory of Auctions and Competitive Bidding," *Econometrica*, September 1982, 50, 1089–1122.
- Mortensen, Dale T.**, "Job Search, the Duration of Unemployment, and the Phillips Curve," *The American Economic Review*, December 1970, 60, 847–62.
- Parker, Robert M., Jr.**, *Bordeaux: The Definitive Guide for the Wines Produced Since 1961*. New York: Simon and Schuster, 1985.
- Riley, John G., and William F. Samuelson**, "Optimal Auctions," *The American Economic Review*, June 1981, 71, 381–92.
- Thiel, Stuart E.**, "Some Evidence on the Winner's Curse," *American Economic Review*, December 1988, 78, 884–95.
- Wilson, Robert**, "A Bidding Model of Perfect Competition," *Review of Economic Studies*, October 1977, 44, 511–18.