Gauging the Art Market
By its
Broad Indicators
And
Never Mind Bill Gates

Readers might remember in the early eighties, when the prevailing interest rate was 18%, it was reported that Dan Terra had paid $3,000,000 privately for Samuel F.B. Morse's **Gallery of the Louvre**.[1] At a time when just a handful of American paintings by some of the heavies had barely topped $1,000,000 this was incredible news.

Trouble was, what we were hearing and reading was less than half the story. Dan wanted to send a signal that he was an important buyer so he broadly announced his acquisition of the Morse always including the price he had paid. For those of us who were accustomed to paying with cash for what we bought, the implication was that Dan, too, had paid cash for the Morse.

However, years later the Cultural Ambassador-at-Large Terra told me he was paying the purchase price in installments over 25 years. Since, in the art world, it has long been customary that notes bear no interest, Syracuse had agreed in effect to forfeit the prevailing interest rate, or 18% on what was owed to them from Dan, which allowed him to discount $3,000,000 to its Present Value, or about $48,000, perhaps the right price for this somewhat esoteric work but in any case a far, far cry from $3 million cash.

I cannot help but wonder if the report in the ARTnewsletter "Bill Gates: Recent Record-Breaking Acquisitions"[2] describes a similar case. According to the report, Gates recently bought two important pictures, a Chase for a reported $20 million and a Hassam for a reported $10 million. The report goes on to say that in the last five years, Gates has bought five American paintings at megawatt prices, two for near $10m each and three for upwards of $20m each, this Hassam, a Homer called **On the Grand Banks** and George Bellows's **Polo Crowd** sold at Sotheby's on December 1, 1999. Wow!

This is interesting, if not salacious news. But let's keep it in perspective; what should one who is not closely engaged in selling paintings to Bill Gates conclude from this? That the whole American market is booming as never before? No one needs to be told again that Bill Gates, who has given American Art "a shot in the arm"[3] according to one informed source, is the world's richest man. But what if the circumstances of these sales are more

---

[1] The writer alleges no information whatever about the transaction other than what he has read in the press and heard on the street.
[2] July 24, 2001
[3] Ibid

like the Terra deal than a straight cash deal, as the report implies? In that case, readers should not infer the same implications about the robust health of the broad market for American paintings as that suggested in the story.

Auction Sale Reporting: The Present System as Statistical Nonsense

Auction results are always reported in the press in a time-honored but folksy manner. Reports following sales at Sotheby's and Christie's invariably start with an absolute grand sum of prices and a recitation of the various record prices achieved in the sale, all of which include premiums. We then learn what percentage of lots was sold of the total number that was offered that day or evening. Then, a percentage of what the absolute grand sum of prices was, again including premiums, of the total estimated value of the sale, but this time without premiums.

For those market watchers who would like to draw conclusions about the growth and direction of the broad market, none of this mumbo-jumbo makes good statistical sense. In fact, anyone *including* Bill Gates who draws conclusions of any sort from this mélange - and plans to act accordingly - is going to make some pretty bad decisions. Obviously, unless you are the sale's only consignor, the total sold including premiums as a percentage of total estimates without premiums for all lots offered is a virtually meaningless figure mainly because it compares apples to oranges.

We'll get to apples and oranges but let's talk for a minute about "Record Prices" and "market growth."

First, if volume by lots is a measure, the auction market at Sotheby's and Christie's, ordinarily steady at between 140/180,000 lots each annually, has not "grown" in ten or more years; in fact, it is shrinking. Second, if gross sales by dollar value are the standard, sales in 2001 are only about 60% of what they were in 1989. And third, but most importantly, no sales figures from Sotheby's or Christie's are ever adjusted for inflation or for the so-called "wealth effect" meaning that "record prices" are a better indicator of some macro-economic index other than increased interest in paintings. Nevertheless, by now it is a fact: the press reports "records" of some description, in up markets and down after every auction sale.

But what is a "record price" for say a deKooning if this year's offering is not the same as last years offering? What does a "record" price for an American work on paper mean if this year the work on offer was a pretty good example by Eakins and last year it was a masterpiece by Homer? These are oranges and apples; in other words we cannot calibrate the market for works of art, which are all unique, without an objective standard to measure one against another.

The customary policy has always been to gauge the whole market by its isolated cases or outliers, the very few notable prices brought, like those in the Dan Terra or Bill Gates stories, under extraordinary or even private circumstances. We are to presume that these

reports tell us something meaningful about the state of the market for other pictures of the type in question and, by extension, for all pictures in the category. When a painting by so-and-so fetches a "record price" of $3,000,000, surpassing the old mark of $2,000,000 set last year, readers are expected to infer 50% upward progress in the art market.

However, for this assumption to have validity, three conditions must be met:

a) The data must refer to fungible or interchangeable assets, i.e. apples and apples;
b) *All* other factors bearing on the market including time, must be equal;
c) All measurements used in the comparison must be calibrated identically.

To meet the first of these conditions, data given as comparable must refer to like entities. Strictly speaking, it is not necessary that the objects of comparison be *virtually* identical only that they be fungible, like a ton of sugar or a regularly scheduled flight from Chicago to New York. A ton of sugar exported from Jamaica and sold on the open market at the Mercantile Exchange in New York is not the very same ton of sugar made in Costa Rica and sold on the Chicago Board of Trade but they are, for all practical purposes, indistinguishable from one another. By contrast, the "product" offered by Sotheby's and Christie's is not fungible like most commodities but limitlessly diverse.[4]

For the second condition to be met, all other factors that bear upon the data must be constant including, but not limited to: time and date, the money supply, interest rates, information symmetry, supply and demand, the Consumer Price Index and real income to name but a few.

The influence of just one of these variables, a change in real income for example, can be enormous. In absolute terms, airplane and theatre tickets, butter and real estate prices are at all-time highs today. For those whose net worth has changed modestly during the boom years of the mid to late nineties, the prices brought by paintings and Fifth Avenue apartments must seem staggering indeed. However, for those whose net worth has grown one hundred or a thousand fold during that time, the price of a Picasso could actually have gone *down* from the "high" established in the late 1980's. To such individuals, art and real estate are 'cheaper,' in real terms, than ever before. From this perspective, record prices for paintings at auction serve more as an indicator of growth in real income than as an indicator of 'growth' in art prices.

Finally, prices even for standardized measurements of fungible commodities must be stated in the same currency. However, the influence of a strong Japanese Yen can bring extraordinary upward pressure into the London art markets against which British or American buyers, when their pounds or dollars are weak, are at a disadvantage. And, since record prices are always calibrated in dollars, those with only dollars in hand receive a distorted picture of the market.

[4] The term "Product," as it is used here, refers to the hard goods transferred through the auction process by either company.

Obviously, the failure to maintain what is known as *Ceterus Paribus* or "all else being equal" is not restricted to art market indexes. People from countries other than the United States, who have Lira, Yen or Pounds Sterling, do buy stock in US markets. The most popular stock market gauges, the Dow and the Standard & Poor's for example, make no pretense to holding all other factors equal and, in fact, as these gauges trend broadly upwards, they reflect growth in other variables as much as in equity valuations. They do not, on the other hand, attempt to compare apples to oranges or kiwis to bananas. This condition then, to compare like to like, is necessary but not sufficient.

In summary, the problem with tracking art prices, reporting sale results and making projections for sales yet to come is simply that the three necessary conditions for making valid and comprehensible comparisons from past market events are absent.

However, notwithstanding that conditions for auction sales are not kept constant, that fine art property is not fungible, and that currency markets shift continuously, market events for *all* property offered at Sotheby's and Christie's have two enduring characteristics in common:

- The pre-sale Low Estimate and
- The Hammer Price.

If data were collected plotting a frequency distribution of hammer prices around a mean of pre-sale estimates, comparisons of sale results for any category of property could be drawn between their respective market events.

Information Technology: A Very Short Primer of Useful Terms

Before we go on, we need to define some of the terms I will be using throughout this paper. If you are with me so far, but were set back some by that last paragraph, this exercise is for you.

In computing descriptive measures in statistics, interest usually focuses on two measures:

(i) The **mean**, a measure of central, or average, value of the data in a series and
(ii) The **variance**, a measure of the degree to which observations are spread out about this average value.

A mean can be **absolute** (fixed), as it is for archers aiming at the bull's eye on a target or **floating** (relative) as it is for pitchers throwing at the strike zone in baseball. In the first instance, each player, no matter in what field or country he plays, aims at a target of identical measurements. In the second, pitchers try to hit a strike zone continuously redefined as the vertical distance between the 'letters' and knees of each individual batter shifts. Notwithstanding, the mean, in either case, represents an **expectation** and throws, or more precisely the spots where throws land, represent **observations** in a frequency

distribution around the mean. Observations may fall to the left, right, above or below the target and the distance to the target from the spot where each throw lands is called a **variance**. A pitcher or archer whose throws, when taken together, fall consistently in, or just on the edges of, the target would have a smaller variance than a pitcher whose throws fall consistently outside the target. A pitcher, who throws consistently either in or just out of the zone, would have a low **volatility**. A pitcher, who throws sometimes in but sometimes way out of the zone, would have a high volatility. For this reason, high volatility is correlated with **risk**.

If it were known that, in the **aggregate**, 70% of all sixth grade graduates in the United States were likely to graduate from high school and that, in the aggregate, 70% of all high school graduates were likely to go on to college, we would still know nothing about little Johnnie's chances of graduating from high school or college. In the example, any sixth grader's chances of graduating from high school and then college were derived from analyzing data comprising the whole **population** "all sixth graders" and "all high school graduates." However, a **sample** of that population, or just a fraction of all those in the population, might have sufficed to yield the same result provided the sample were not either internally or externally **biased** or self-selected. The chances that 7:10 sixth graders will graduate high school can be extended to any group of sixth graders but not to any sixth grader in particular.

These terms will occur in the discussion that follows.

Knowledge Management: very lately applied to the Auction Industry

The computer age has been slow to take hold in the art industry. As late as 1995, experts at Sotheby's had no computer terminals on their desks and no intranet for sending or receiving simple memoranda. While Christie's has had an intranet since the early 1990's, it was, until the move to Rockefeller Center, a DOS driven platform almost useless for any application beside email. Even today, computers in the expert departments at both houses are used mainly for word processing, as if the machines were high-speed typewriters, nothing more.

Information Technology, the systematic collection of data, and Knowledge Management, that which is learned from data once collected, are not terms spoken at either firm. It is odd, when technology for sending images instantly and cheaply to anywhere in the world has been available for ten or more years, that both companies have relied until very recently and still do rely upon film photography and Federal Express to get images to their own offices in distant countries. Images are, after all, what these firms traffic in and images are reducible to bytes.

Reasons for this technological backwardness in the industry are complex. For one, insiders, at both Sotheby's and Christie's, have forever taken refuge in the knowledge that their business is arcane, not subject to the mechanisms of academic finance nor even to the boilerplate laws of supply and demand. It is perhaps that fine art property is only

subjectively appraised and, therefore, difficult to quantify that the industry has never been subjected to the kind of rigorous, fact-grinding analysis that is everyday presented in the daily papers covering the financial markets.[5]

Notwithstanding, it is a fact that decisions of staggering financial consequence are daily made on the barest supporting data at Sotheby's, Christie's and now Phillips.[6] Since, at these firms, new business has always come streaming in the front door, there has been little effort, and fewer resources, devoted to data retrieval systems.

It is not that experts fly entirely blind. Expert departments at both Sotheby's and Christie's collect sale data covering the period from roughly 1970 to the present. Records are crudely collected and more crudely kept in a manner consistent with how they are intended to be used: college-aged 'interns' with scissors and glue pots disassemble past sale catalogues, removing descriptions and illustrations. Then, the interns paste the information onto 8 x 10 inch shirt cardboards. Without knowing what or why they are writing, the interns annotate the cards with prices and other information and file them alphabetically by artist's name in banks of upright file drawers. The system, which has its roots in an Industrial Revolution-era paradigm, takes up vast amounts of valuable real estate and is fraught with all the dangers that loss, misplacement, error or omission can devise.

The system is useful, after a fashion and of course many experts with decades in service have impressive, deep memories. Yet memories are not infallible and they, too, were designed to function, like the card systems, on an essentially flat platform, *i.e*. they recall data in two dimensions, price and date, one record at a time, not relationally and never aggregately.

In addition to the in-house system, experts can rely upon a few other aids. Printed and bound indexes of prices at auction have been available for twenty-five years or more. But they are not illustrated and they are organized by decades and otherwise not indexed. On-line references for these data, such as Artnet.com and ArtPrice.com, only four to six years old, are an improvement over the books and they have reduced dependence upon, but not made the in-house system obsolete.[7] But there are problems: the on-line databases are constructed as a complete population of sale results collected from even the most obscure auction facilities. As such, they comprise data that are unscreened for important

---

[5] Frank K. Reilly, of the University of Notre Dame, calculated "Risk and Return on Art and Antiques" (unpublished mss. July 1992) cited in Frank K. Reilly and Keith C. Brown, "Investment Analysis and Portfolio Management," (New York: The Dryden Press, Harcourt Brace College Publishers, 1997). The writer of these pages has made several attempts to secure copies of these studies, thus far unsuccessful. It is not known, therefore, what data Dr. Reilly used for his study nor what constituted a "risk" or a "rate of return." It is likely that Dr. Reilly used the "Sotheby's Index" published as a joint venture between Sotheby's and Barron's Weekly. In this series, prices for various classes of property were tracked and averaged as evidence supporting a rise or fall in levels for the underlying classes, relative to performance in a benchmark year.

[6] "Smooke Collection awarded to Phillips, de Pury & Luxembourg for fall New York auction" in The New York Times, July 26, 2001

[7] Hans Neuendorf, President & CEO of ArtNet.com tells me that experts at Sotheby's and Christie's are his biggest clients in the retrieval of their own data.

considerations like authenticity, making them very low-grade resources. They also comprise variables that are qualitative (nominal) like media and dates of execution which, because they are not quantitative (interval) permit only ranking or counting, not calculations.

These services, available by subscription, are costly and have a number of other notable shortcomings, among them:

- Databases are time-series ordered, not cross-sectional
- Users may sort and filter but cannot manipulate or query the data in accordance with criteria of their own design
- Data can only be referenced in single observations, like a book, not as aggregates
- There is no, or quite useless, data analysis
- There are no credible formulae for discerning trends

A new formulation for art market data analysis is lacking and sorely needed.

My Database:
How collected, Source, Size, Breadth, Prevalent Bias and other Notes

Source

My descriptive data are collected, as are all others, record-by-record, from catalogues published by Christie's and Sotheby's and entered manually, into Access 2000. Illustrations from the catalogues are scanned, adjusted and converted to jpeg files, batched by sale and lot number and saved in a separate *.mdb file apart from the data to which they are linked. Recorded prices are collected from the post-sale price lists disseminated by each house after their respective sales.

Size and Breadth

Data collection at ArtPocket.com is confined, at this point, to sales at Sotheby's, Phillips and Christie's American painting departments between 1983 to the present. I have included only major sales held at these houses in order to capture a data set from the top end of the price spectrum. Only catalogues appearing in March, April/May, September and November/December were scanned.

Bias

Inasmuch as consignors of low-end market property are not apt to practice risk management, I have not included these data in my records. Therefore, I make no attempt to present a definitive record of *all* sales of American paintings offered at either Christie’s or Sotheby’s over the period. As they are virtually unlimited, paintings appearing in minor sales, at Christie’s East or at Sotheby’s so-called Arcade sales,

Americana and Decorative sales are likewise not included. The population of market events is already dramatically right skewed without the low-end market segments, more so with them.

It is important to recognize that the ArtPocket database is not a compendium of artifacts to be referenced like a book but a set of market events involving those artifacts. A single work of art, described identically in each appearance, may be entered once or numerous times to record separate market events on different dates or at different places. Inconsistencies in a painting's variable characteristics such as title, medium, size and date of execution are therefore immaterial.

Common-sizing: Bench-marking sales at auction

To make valid comparisons between one market and another, Statistical Analysis needs a common reference; something every item compared has in common. But because paintings are unique, not fungible, a reference point or benchmark measuring sales of different works has never been adequately identified. Therefore, this study treats each sale lot not as a unique work of art, which each lot undeniably is, but as a unique market event, something every auction lot, no matter what it is, undergoes like every other.

Baseball players are similarly unique, one from the other, but each player participates in an at bat event where, like every other player, he faces a pitcher and eight defensive players. The manager of the team then in the field expects a certain performance from each new batter based upon past performance and adjusts his defense accordingly. Art experts expect a certain performance from each lot in their sale also based upon past performance and they adjust their estimates accordingly. This is an important concept and it may require a little more illumination.

In 1930, a newspaperman asked Babe Ruth what he thought about being paid $80,000 a year, almost double what Herbert Hoover made. Ruth's response, that he had had "a better year than Hoover" is interesting not because he ignored the vast difference between *what* he and the president of the United States were paid to do but because he inferred that the disparity in their respective compensation packages had to do with how *well* each of them did relative to what they were *expected* to do. This point is central to my theme.

All markets change with time, as the different variables affecting sales rise or fall. However, most market indexes, as for corn or iron ore, depend upon constants, or factors that do not vary. For example, a ton of wheat harvested in Nebraska is for all intents and purposes identical to a ton of wheat harvested in Iowa. Because of this fact, the price of any ton of wheat can be calibrated as an absolute, or in varying dollars per constant ton. This system permits geographically diverse buyers and sellers of wheat to make short-term decisions regarding the price of this commodity in relationship to expectations of what the price will be in so many months. The expectation, while stated as an absolute today, is relative to the market price prevailing yesterday or last year, either higher or lower. A speculator's estimate of the future value of the commodity is based upon

information he has that others may not have or may interpret differently, leading them to different expectations.
It is for this reason that the data I collected for this study comprise a sample (some records) not a population (all the records). The database is not, therefore, a comprehensive record of everything sold at auction by a particular artist in a particular season to be referenced one record at a time but a sample of that population, to be considered together. As such, they may be enlisted to demonstrate the difference between expectations (estimates) and results (Hammer Prices) of American paintings either alone or in the aggregate. They are analogous to the "earnings surprise" on Wall Street, which compares actual earnings to what had been forecast by analysts for the quarter.

Here, the measure of an earnings surprise for one lot at auction is stated as a Variance – how far the Hammer Price is from the Low Estimate - from the mean of observations for that artist. Deviation from the mean, or in this case, variance, is a measure of uncertainty. The measure of a painting's absolute value, which it ostensibly has at all times, is of less interest than its value relative to its expected return. This is because understanding and managing risk, which a picture does not acquire until it is consigned and which it discharges when it is sold, is how punters measure the value of their own assets and govern their future market activity. How do we compare a Cassatt painted in 1868 and sold in 1998 to a Bierstadt painted in 1888 and offered in 2000? By basing the comparison not upon absolute value but upon performance relative to expectations.

Estimates – a History and an Explication

In financial circles, the value of an asset is the present value of its expected return. Specifically, one must discount an expected cash flow by the required rate of return applicable to this asset class. Estimates for fine art property to be sold at auction are, however, derived somewhat differently.

In the auction industry, estimates are based upon a loosely defined 'formula' held in each expert's head. The formula takes into account artist, subject, size, quality, date of execution, condition and market precedents. Three to four months in advance of sale, the expert does what an expert does and he or she projects a present value for the lot ordinarily not discounted by the time interval still to pass between the time of consignment and the date of the auction.

Experts in the painting department at Parke-Bernet, where I worked from 1967 to 1974, always assigned estimates, sometimes in concurrence with the consignors' expectations but sometimes without. Estimates were assigned well in advance of sale date, sometimes four or more months in advance. They were determined as they are today, relative to known values achieved for similar examples from the recent past, sometimes a year or even two years past, plus or minus an adjustment for subject, size, date, condition or 'quality'.

Consequently, estimates were and continue to be, by definition, essentially backward-looking. They are often set with reference to 'hard records' but are most often divined off the top of the head from what the expert 'knows.' In spite of the fact that the sale date for a particular painting was scheduled at some distant date, no allowance for time to market was, or even now is, factored into the estimate.

Estimates were never shared with the public until about 1968 when their existence suddenly became known. Before then, estimates like reserves today, were traditionally confidential between the auction house and its consignors and were not published. Consequently, the confidentiality of estimates and reserves was fiercely defended as a proprietary right of the house. However, when their existence became known, the public quickly fixed upon them as a guide for what to expect in the salesroom on the day of the sale and the policy to keep them secret came under attack.

For a brief period, between 1969 to 1970, and in response to growing demand, the indefatigable Elizabeth Duckworth read estimates on the exhibition floor from a book held very close to her chest and limited to three or four to a customer. This system was quickly overwhelmed. So eventually, but reluctantly, and to Mrs. Duckworth's eternal gratitude, estimates were printed on sheets and posted in the salesroom.

Still later, estimates were printed and, again very reluctantly, included with, but not bound in the catalogues. Before long, demand overtook this policy too so that estimate schedules were bound into the catalogues - but in the back, so that they would not be too readily associated with the individual lots to which they referred. Finally, in 1973, to everyone's relief, management capitulated entirely; estimates were printed in the catalogues alongside their respective lots. This custom continues today.

Viewing this drama in retrospect, it is difficult to recall just what all the fuss over releasing estimates was really about. For one thing, an auction houses' customers were then thought to be its consignors; buyers were taken completely for granted. Notwithstanding, the controversy was also about information symmetry; giving more information to the buyer meant shifting more power to him. Today, buyers are held in slightly higher regard and estimates are an indispensable part of the auction process but they are widely misinterpreted and misused.

Estimates are not so much a guide to "value" as they are a measure of the seller's inclination or disinclination to let his or her property go. A reluctant or ambitious seller will demand an aggressive estimate and, if the picture is desirable, the house may agree; after all, the "real" value of the property is yet to be determined. But potential buyers will judge the over-estimated lot as relatively unavailable to the market. Consequently, demand is constricted and the lot's chances of selling will be greatly reduced.

Conversely, estate property or "sleepers" can be transparently underestimated. Astute buyers will classify these lots as readily available. In either of these cases, the published estimates are an unreliable reference for making accurate predictions of "value" and

prospective buyers are thrown back in these instances - if they could discern which instances these were - upon their own resources.

To auction house experts, estimates are actually marketing tools rather than predictors or indicators of "value." This skews estimates in ways not discernible to any but the most sophisticated buyers. Wide spreads and inflated high-end estimates can capture the imagination of ambitious consignors predisposed to hoping for pie-in-the-sky outcomes. But consignors and buyers should focus on the low estimate, the business end of the spread. Low estimates can be identical with, or quite close to the critical threshold at which a painting sells or does not sell (the 'reserve') while the high estimate carries no significance of any consequence, beyond its appeal to greedy sellers.[8]

Estimates take on all these complex characteristics and more. Yet, on balance, they are as good a predictor of "value" as any. Consequently, I recognize them, rightly or wrongly, with all the importance and significance with which the public and the auction house invest them. Furthermore, estimates assign the best available expected value to a property for the four months prior to sale date.

The Buyers' Premium

On May 30, 1975 Christies radically changed its commission structures by introducing the Buyers' Premium. Heretofore, consignors had paid all auction house fees, which, from earliest time, slid on a scale ranging from 25% for lots selling for less than $1,000 down to 1-2% for paintings selling for over $100,000.[9] Sotheby's, adopting the identical concept, followed suit the following Monday June 2, 1975. The premium was charged for sales beginning with the new season in September 1975.[10]

Management at both houses knew that the competitive arena of the business was defined by sellers, not buyers, so that requiring buyers to pay a premium would not detract from business in the same way that raising fees for sellers would have. Consequently, a 10% surcharge was henceforth to be applied to the hammer price of all articles sold at auction and this charge was payable by all buyers, whether private, dealers or institutional.

On January 1, 1993 an additional surcharge of 15% was assessed to buyer accounts on the first $50,000 of any successful bid for property bought at auction at Sotheby's and then the customary 10% was assessed on the balance of the hammer price. Christie's, introducing the identical concept, followed suit in March. [11]

---

[8] In paintings sales at Sotheby's and Christie's, reserves can be equal to, but never in excess of, printed Low Estimates. As a practical matter, 100% of lots appearing in the catalogues are "protected" by reserves. Approximately 50% of lots have reserves equal to Low Estimates.

[9] Commission rates as published were seldom if ever paid in practice. Zero rate commissions were easily negotiated at either auction house for any property valued at above $50,000.

[10] At this time, Christies was a public company. Sotheby's did not go public until June 1977.

[11] Christie's announced a schedule of fixed sellers' commissions on March 9, 1995. Sotheby's, again with the identical concept, followed on April 13, 1995.

It is impossible to assess the impact of this change upon sale prices. While the new commission structure was introduced amidst howls of protest and predictions that the entire system would collapse, it is now generally thought that prices were never materially affected by the extra weight of the charges. It is interesting to note, however, that dealers, who were particularly incensed, bullied the auction houses to henceforth report to the press prices that included the premium, as if this would somehow embarrass them into dropping it. The effect, however, was quite the opposite since it played nicely into the strategy to report "record" prices at ever higher and higher levels.

Since its inception, the buyers' premium has frequently changed and since that time, it has been different at different times at different houses. Notwithstanding, consistency in all variables is important when comparing observations in a data set. Therefore, in order to compare auction Prices this year to auction Prices in any previous year, Buyer's Premiums need also to be consistent. And, because my system compares results to expectations and because estimates *never* impound premiums, all Buyers' Premiums have been backed out of this study to give consistency to the results from the beginning of published estimates at Sotheby's, *circa* 1971-72, to the present. This means that I have backed out the Buyer's Premiums from all recorded prices issued by the auction houses that include it, returning the hammer price once again to its venerable position as a standard.

OK, School's Out!

Even with the old reporting system, most American art market watchers know intuitively that there is a difference between Sotheby's, Christie's and Phillips American Paintings Departments. Broadly and succinctly speaking, it has seemed for the past twenty-odd years that Sotheby's always has had smaller, better 'quality' sales, fewer Bi's, higher prices - especially at the top end - and a consistently bigger sale's gross. Phillips has had two remarkably bad American sales in its short, new life in the big-time. But these are subjective statements; now that we have the data, let's see how they really stack up.

When one charts the numbers in a data set, patterns emerge. Imagine that you are a new consignor. You have never been to an auction and you are entering the market for the first and only time in your life with grandmother's Cassatt. The expert at Sotheby's tells you it is worth $1/1,500,000, the expert at Christie's tells you it is worth $2/3,000,000, and at Phillips they say $3/4,000,000. The experts all roll out "statistics" intending to demonstrate how many paintings like yours they have sold previously. This scenario, totally fictitious, is not at all unusual. Nonetheless, there is a broad discrepancy in these three assessments and while any dope would rather sell for more, you wonder: could making estimates be that inexact? What are you supposed to do now?

The answer lies not in knowing how many records one house or the other has established for Cassatt but in how closely the expert you are talking to has hewn to the target expectations he or she has historically set. This single number can be extended to your consignment as a fair approximation of what you may expect.

The data tell us that Christie's has demonstrated a variance for Cassatt of 0.5382 while at Sotheby's they have demonstrated a variance for Cassatt of just 0.2924. These numbers, with a range of between 0.00 to about 0.60, are complex, not intuitive. Never mind: all you need to know is that higher variances indicate higher deviation from the mean, which implies that where Cassatt is concerned, Sotheby's has hit its mean with far greater reliability than Christie's no matter what the painting and no matter what the estimates. In plain English, the data indicate that on sale day Sotheby's has more regularly delivered on or near expectations, yours and theirs, while Christie's has more regularly deviated widely from expectations.[12] Since a higher variance implies greater risk, when the novice consignor has Cassatt to consign, he or she had better go to Sotheby's.[13]

More telling, between 1983 and the present, Sotheby's has demonstrated a variance overall of 0.1205 while Christie's has demonstrated a variance of 0.1546. These figures aggregate the records for all artists and all Hammer Prices relative to all estimates over the entire period 1983-2001. While the difference may appear small, Christie's variance is nearly 30% higher on its face than Sotheby's.

The following charts display other patterns for the category American paintings that I think readers will find interesting. Some, like Bought-in percentages against numbers offered, are familiar. But seeing it as a continuum over many years or even decades is not.

[12] Astute readers would be correct to note that large variances include Hammer Prices considerably above estimates in with those considerably below. Inasmuch as Hammer Price distributions are constructed in a linear fashion, not scattered, a higher volatility could indicate large, positive earnings surprises. However, higher volatilities always indicate results that are wide of the expectation implying that estimates were off one way or the other. Therefore higher volatilities still indicate higher risk since in any business, prices that come in closer to expectations are more desirable than lucky breaks.

[13] There are not sufficient data to compute the variance for Cassatt at Phillips.

Exhibit 1 – Bi Rate by House and Year

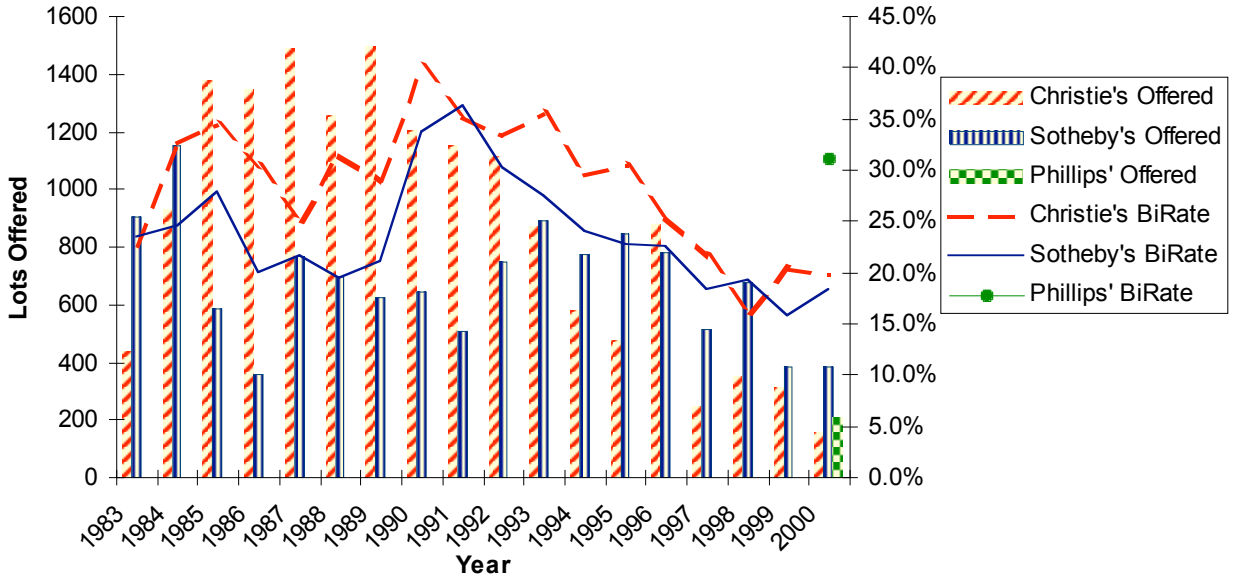

This chart is drawn from merely counting the observations in the published sale data spanning the period 1983 to the present. In the bar graph, we see volume at Sotheby's and Christie's American Paintings through the years 1983-2000. Notice that in the market run-up that took place between 1986 through 1990, Christie's American Paintings consistently offered near twice as many lots as Sotheby's, which appears to have had a less well-defined policy. We also see (in the line graph) that a wide disparity in Bi percentages at Sotheby's and Christie's has existed for near two decades almost without interruption describing a roughly parallel trend and a perennial gap. Interestingly, the lines cross momentarily at cycle beginnings/ends, implying that Sotheby's has been somewhat slower to perceive changes in the business climate than has Christie's.[14] The chart also implies that until 2000, the Estimate/Reserve Paradigm, a staple of the business plan at both houses, has been more harmful to Christie's and its American paintings' clients than to Sotheby's and its clients. In fact, Sotheby's has enjoyed the broad outlines of a Competitive Advantage in maintaining the paradigm. We also see that in 2001, Bi percentages have risen above record low levels set in 1998, an indication that the cycle has peaked.

[14] Readers should note that the administration of Sotheby's American painting's department has been consistently, or near consistently the same for twenty-odd years while in that period Christie's American department has changed heads no less than five times.

Exhibit 2 – Bi Rate by Lot Bands

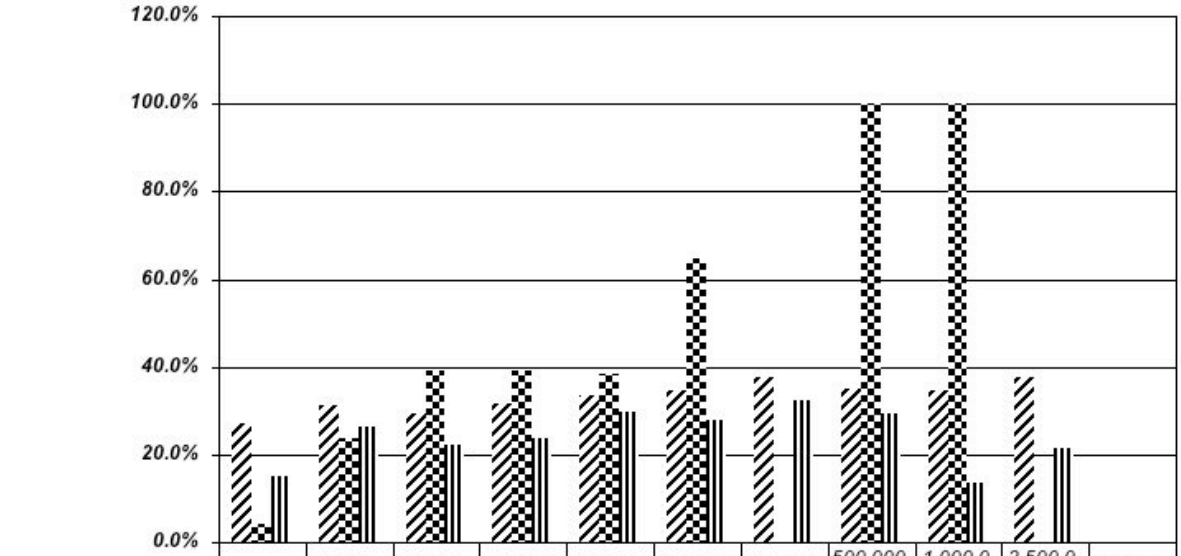

	0-2,000	2,000-7,500	7,500-30,000	30,000-60,000	60,000-100,000	100,000-250,000	250,000-500,000	500,000-1,000,000	1,000,000-2,500,0	2,500,000-5,000,0	5,000,000-
Christie's BiRate	26.9%	31.2%	29.5%	31.6%	33.4%	34.5%	37.7%	34.9%	34.8%	37.5%	0.0%
Phillips' BiRate	4.0%	23.9%	38.9%	39.3%	38.5%	64.7%	0.0%	100.0%	100.0%	0.0%	0.0%
Sotheby's BiRate	15.0%	26.2%	22.3%	23.6%	29.7%	27.9%	32.2%	29.4%	13.6%	21.4%	0.0%

When Bi's are viewed as a percentage of offerings in the different lot bands, Sotheby's demonstrates a relatively tight rate, ranging between13-32%, but lowest at the top end, where it matters most. Christie's demonstrates a broadly higher rate, ranging between 26-37% with their lowest rate in the bottom end, where it matters least. Moreover, starting with lots priced above $7,500, Christie's generally buys in a greater and greater percentage of its offerings as lot values increase while Sotheby's buys in at a lower rate as lots increase in value. Until the failure of John Singer Sargent's Girl Fishing At San Vigilo (Sotheby's Sale 7668 May 24, 2001) to bring its $4/6,000,000 estimate, Sotheby's record in the over $4 million band had been flawless.

Phillips, which has had only two sales, has nonetheless bought-in 100% of the few lots it has offered above $250,000 but just 5% of the lots it has offered in the under $2,000 band. This would not appear to be consistent with what we understand to be their business model.

Exhibit 3 – Lots Sold % Below Low Estimates

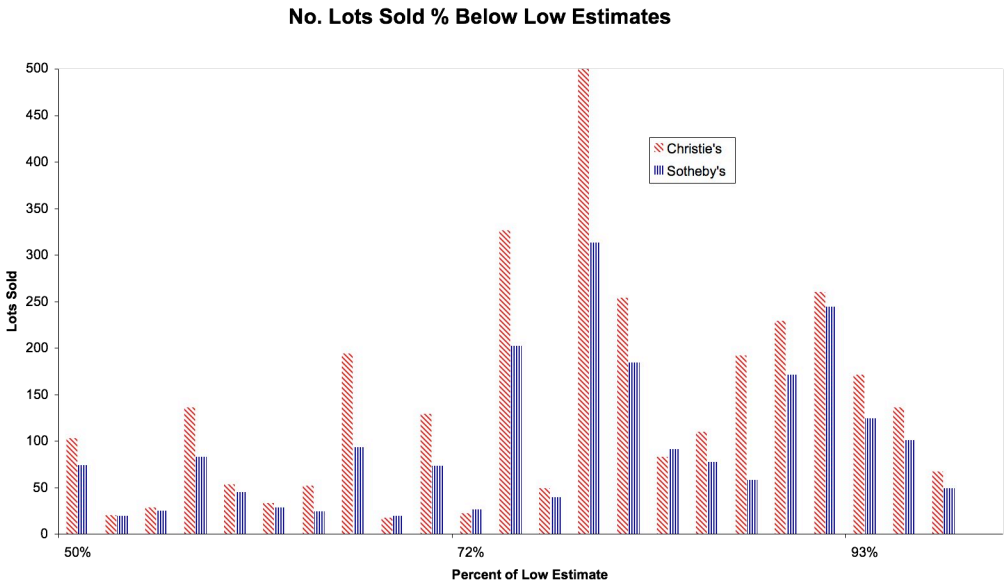

This chart shows how many lots (y axis) both houses sell below low estimates, which we will call Negative Earnings Surprise, and what percentage of all offerings they sell in that range (x axis). The chart indicates that both houses sell most often at the 80% of LE point. However, Christie's American Paintings has sold overall near twice as many lots below estimates as Sotheby's, mostly in the 80% of low estimate range. Curiously, Sotheby's sells below low estimates more often in the 97-98% range than does Christie's implying, I suppose, that once he has passed the reserve, the auctioneer is in a hurry to proceed, perhaps an unconscious form of efficiency. However, since low estimates are a threshold for customer expectations and customer "value," this simple count implies that in selling twice as often below estimates as Sotheby's, Christie's has twice as many unsatisfied customers as Sotheby's.

Exhibit 4 – Variances for Individual Artists (all years)

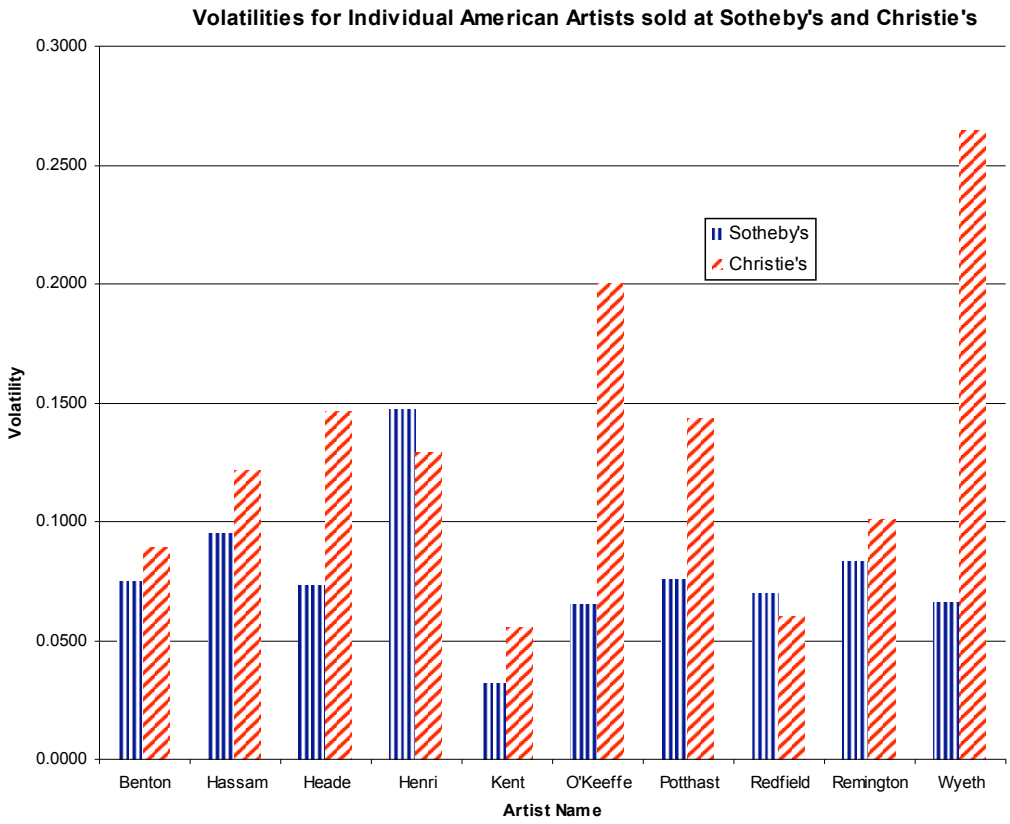

With the volatility for an artist known, one is able to separate risk into its two component parts. The first is Unique Risk, which measures how prices for an individual artist rise or fall irrespective of prices recorded on the broad market; the second is Systematic Risk, which is applicable to all art prices in accordance with the phase of the business cycle or other macroeconomic phenomena. It is important to realize that prices for Bierstadt or Cassatt (unique risk) may be rising as prices for the broad market of all American paintings (systematic risk) is falling.

ArtPocket has analyzed data from thousands of American auction sales from 1983 to the present so that users can derive a Risk Profile for any artist. The variance from its own mean exhibited for works by any artist offered and sold at each house is a measure of risk. Since individual artists exhibit discreet, measurable variances, past performance is a reliable indicator of future risk at each house.
Here I have selected at random demonstrated volatilities for ten artists at each house. In keeping with aggregated variances, Christie's demonstrates higher volatilities for all but two, Robert Henri and Edward W. Redfield. In the case of Andrew Wyeth, sellers who have chosen to offer at Christie's have faced extreme risk.

Exhibit 5 – House Variance by Year

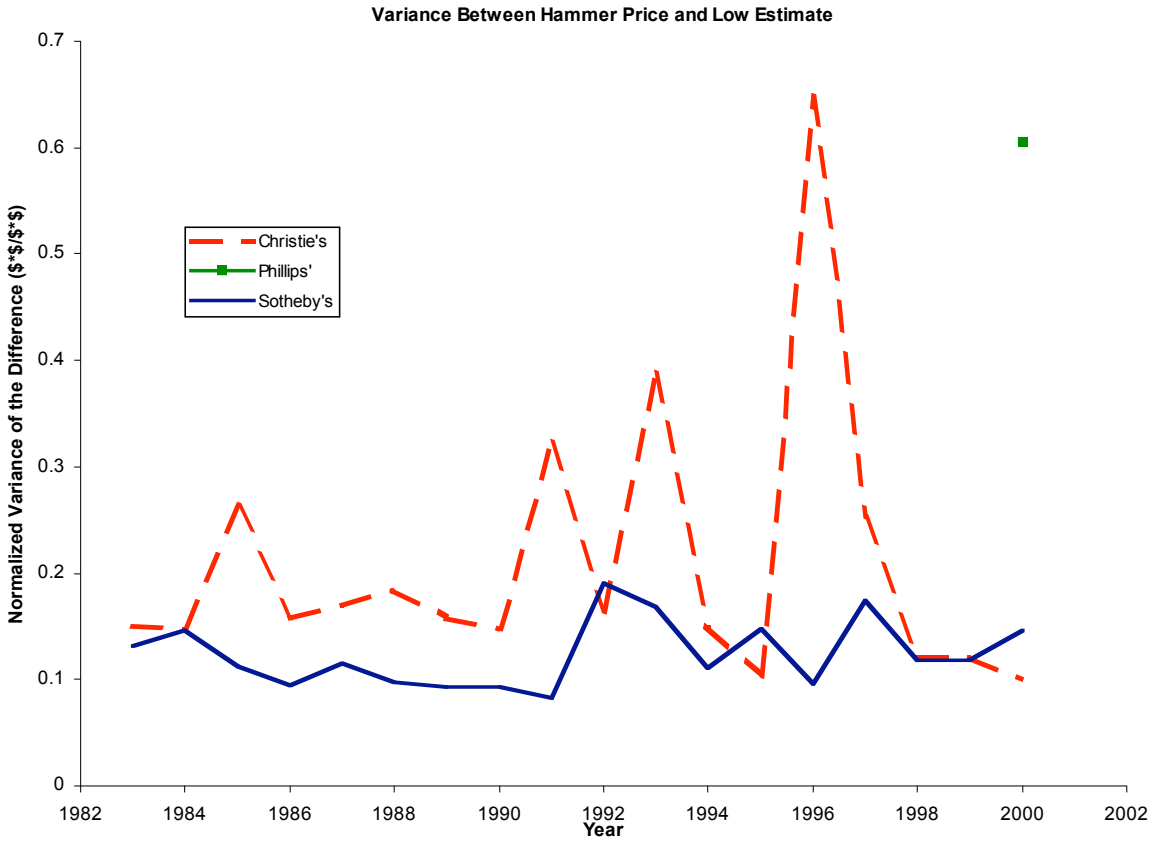

An analysis of the data describing the performance of all American Paintings at auction at Sotheby's in the year or in a set of years returns a different figure than the performance of all American Paintings at Christie's for the same period. This chart measures variances showing how far all Hammer Prices fall from their respective Low Estimates at each house each year. Remember; while paintings are different, variances are comparable because each department sets its own mean. Notice that Sotheby's demonstrates a consistent overall variance from year to year, rising only in the years 1984 and 1992, both cycle ends. Christie's, on the other hand, has demonstrated a sporadic performance over the period covered, rising in 1991, 1993 and 1996 while the market was generally strong. Changes in the administration of the American paintings department at Christie's may explain these spikes.

Charting the Systematic Risk of the art market is tricky business. Simplistically, the computer can analyze a data set for art prices against one or more data sets looking for correlatives among other macroeconomic factors that are largely external to the market for paintings at auction. These might be the indexes of leading or trailing economic indicators, housing starts, interest rate spreads, the purchasing managers' index, or sales

of durable goods. This kind of analysis, called a regression, is beyond the scope of this paper.

We can, however, for these purposes, chart the internal symptoms of Systematic Risk by assuming that in any year, property comes into the auction market and sells in direct proportion to how consignors and buyers judge such external factors for themselves. In that case, the composite weight of art sales in the category might be the best indicator of the market's apparent, systematic risk.

Exhibit 6 – Sum of All Hammer Prices vs. All Low Estimates by Year

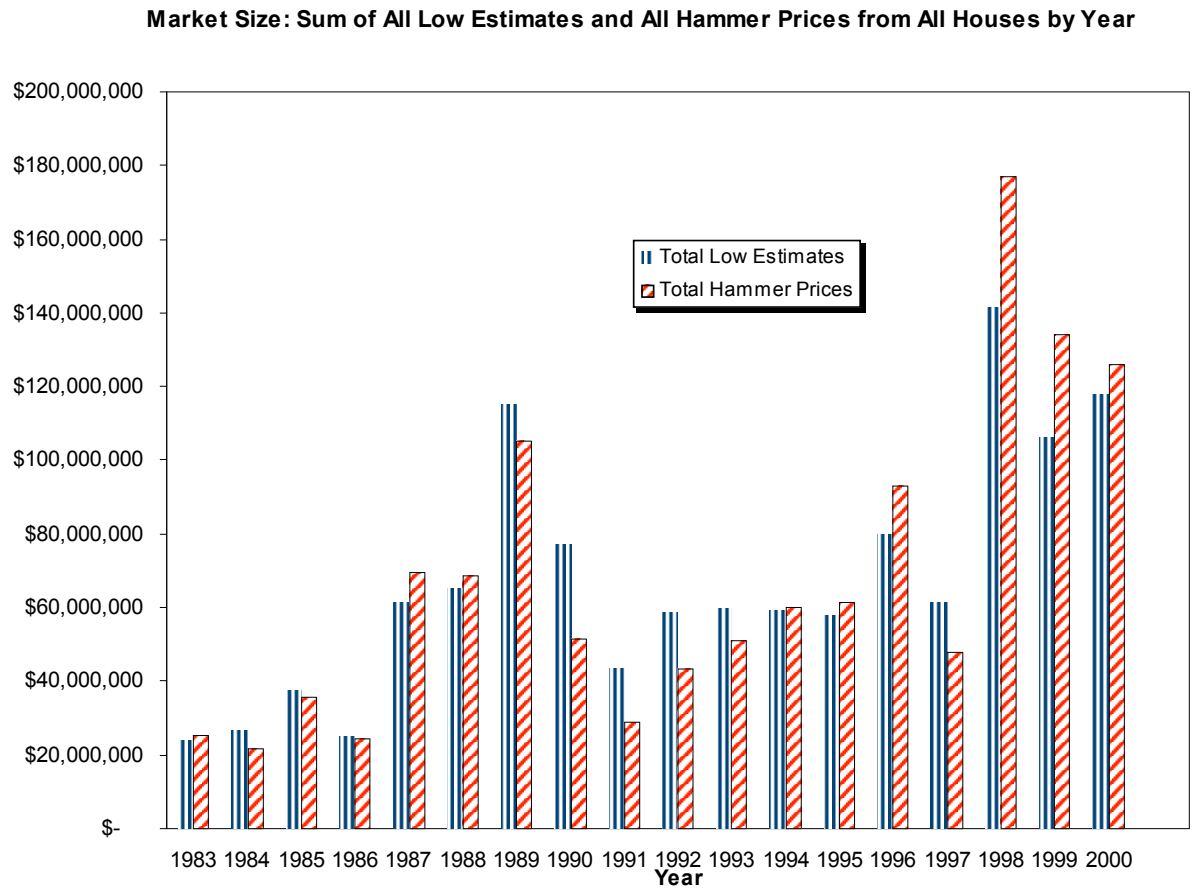

In the chart above, I have summed the Low Estimates of all property coming into the market at Sotheby's, Christie's (and Phillips in 2000 and 2001) by year and charted that figure against the sum of all Hammer Prices for the same property in the same year. These lines roughly mirror the intuitive business cycle, steadily rising until 1989, when the macroeconomy entered a recession and the art market fell off the table. Notably, estimates and prices rise again into 1998 and then fall off sharply into 2000.[15] More

[15] The chart does not include figures for 2001, as it is not yet complete. However, total hammer prices of $51 million without the November/December sale would appear to be well off the pace of $177 million set in 1998.

importantly, aggregate Estimates trail Hammer Prices in the years leading up to peaks and surpass them in years leading to downturns. This implies that, in the aggregate, volatilities are lowest in mid-cycle (estimates are more reliable) than they are either before or after cycle highs and lows. Over the past four (1966-'74; '74-'82; '82-'90; '90-'98), the cycle has demonstrated duration of 8.0 years each.

Exhibit 7 – A/B Line

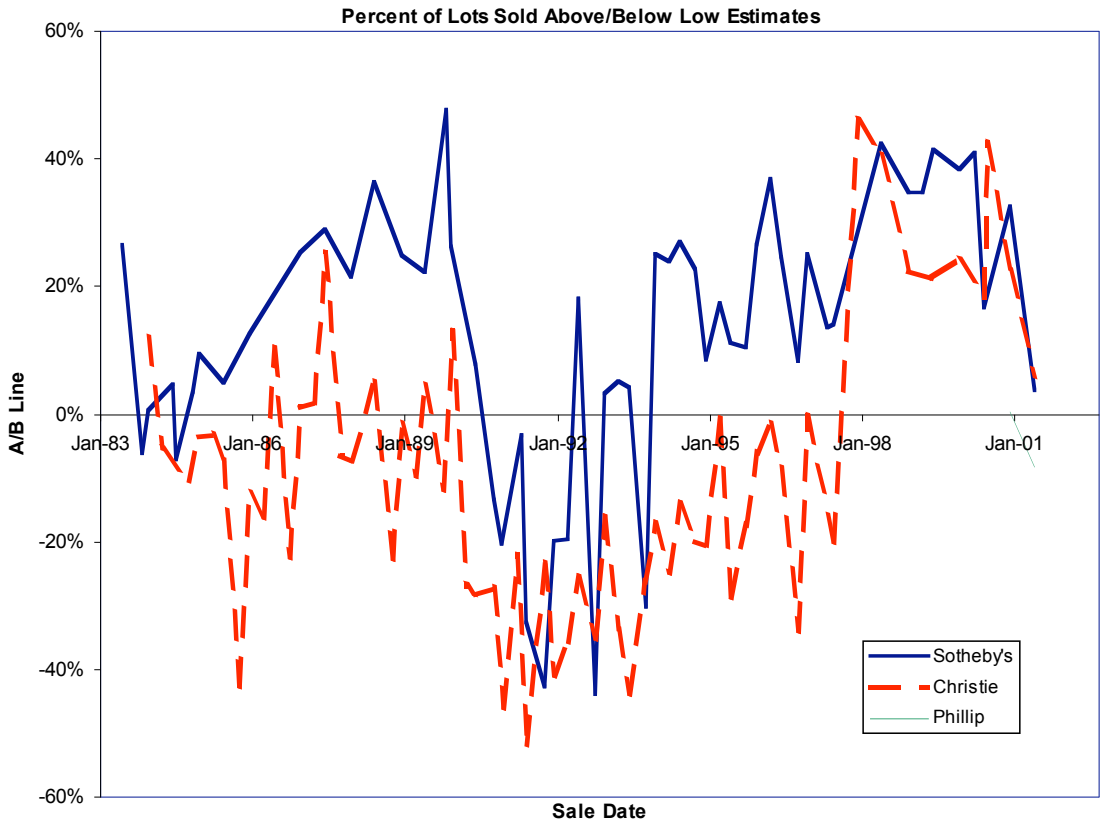

In this chart, we can see a more specific application of the data to the determination of the art market's Systematic Risk. This function, which is analogous to the A/D Line graph published daily in the Wall Street Journal, plots the number of equity issues advancing on the day subtracted from the number of issues declining.

Here the function is defined as the number of lots in an auction sale that bring prices above their own low estimates subtracted from the number of lots that fail to bring prices above their own low estimates. Notably, the Bi lots are grouped in, by definition, with those failing to bring prices above low estimates.[16] Positive results indicate a sale in

[16] In order to equate sales of different size, the data are normalized: the result, above, is divided by the total number of lots offered in the sale, which yields a percentage, then multiplied by 100 to get rid of the decimal.

which there were more lots sold above low estimates than below; negative, a sale in which fewer lots sold above low estimates than below.

Notice that the chart again mirrors the intuitive business cycle; a peak in the year 1989, a drop in 1990-92, the steady rise to 1998 and the drop into the present market. Contrary to popular belief, the jagged, up down pattern in the lines indicates that off-season *i.e.* March and September sales are less risky for consignors than so-called important sales, those held in May and November, where consignors all insist their property be offered. Also notice that Christie's, until most recently, has quite continuously bought-in and/or sold more lots below low estimates than Sotheby's.

Conclusion:

It is important to realize that these charts are not constructed by counting the number of lots merely sold and then stating the total as a percentage of lots originally in the sale, as is the custom. It is important to *sell* a lot, yes, but more important to know by what margin of error.
My system rates painters, experts and auction houses not in absolute values but as Babe Ruth rated himself: relative to how well he was expected to do when he came to bat. I therefore rate art market events relative to low estimates, the lot's expected return and the "business end" of the spread, which almost certainly impounds threshold reserves. And premiums are never reported in sale results since estimates never impound premiums and premiums change. I ignore high estimates altogether. I have looked at the auction business in the aggregate, not as a succession of unrelated, unique events where each lot faces an unquantifiable risk governed largely by chance.

§