Analysis of Auction Price Risk: An Empirical Study of the Australian Aboriginal Art Market[1]

Ilya Voytov

April 14, 2006

Honors thesis submitted in partial fulfillment of the requirements for Graduation with Distinction in Economics in Trinity College of Duke University

[1]I am especially grateful to Professor Neil De Marchi, my advisor, for his guidance and encouragement. Special thanks go to Lynh Le for invaluable help and illuminating criticism and support.

Abstract

Auction theory economists have shown that auctions can be structured to maximize the expected revenue to the seller. In this thesis, I show that they can also be optimized to minimize the sellers' risk through an understanding of the driving factors behind seller's auction price risk. I derive a general form equation for auction price variance, and discuss how changes in the number of bidders and the type of bidders affect the sellers' auction risk. An empirical component of this paper takes data from auction sales of Australian Aboriginal art and uses observed price variance to make deductions about the underlying types of participating bidders.

1 Introduction

The use of auction sales has been an increasing trend in recent years. Businesses are increasingly likely to choose an auction to sell a division or stock, consumers buy and sell goods at online auctions, and governments see public sales as a way to raise additional revenue.

Since the original foundation of theoretical papers by Vickrey in the 1960's, the focus of research has been on designing optimal auctions and understanding bidder strategies [Vickrey, 1961]. Much work has gone into generalizing assumptions to facilitate the understanding of the many new auction formats being put to use in the modern economy. We now understand the full complexity of optimizing multi-unit auctions, and how affiliating bidders' valuations with the information in their possession increases revenue to the seller. At the same time, the assumption of risk-neutrality by the buyers has been relaxed, and is no longer required to achieve the main conclusions of auction theory.

The seller, however, is also exposed to risk as a result of an auction. Auctions are usually held in markets where sales are infrequent, and the usual process of price-discovery through guess and test until demand clears all available supply is not feasible (for instance, when items are so unique that only one is available for sale). Once a seller commits to holding an auction sale, he/she is not at liberty to pull out of the sale once bidding clears the reserve price on the lot. Even if the seller may at that point be satisfied with the revenue from the sale, he/she is still subject to greater uncertainty compared to other more liquid transactions. This may present a challenge to sellers subject to a significant liquidity constraint, such as companies entering bankruptcy, beneficiaries of a deceased collector, or partners in a divorce. Minimizing revenue risk then is a relevant component of auction theory. It is also not an intuitive or a correct conclusion that the same strategies that maximize revenue also minimize sellers' price risk. To take an elementary example, the proof of the equation for the optimal reserve price involves the assumption that the seller would be willing to undertake the small risk of losing a sale in return for a chance to push bids upwards in cases when very few bidders are participating. Here, an optimal reserve price is a condition of revenue equivalence and of an optimal auction, but is directly at odds with the condition of minimizing risk.

One of the booming auction-driven markets is sales of fine art. Recent years have seen evidence for rising valuations of art compared to other types of assets. One peculiar property of fine art, however, is that people tend to have strong opinions on the worth of individual pieces, and since quality is not objective, these opinions vary from person to person. The effect is that, while artwork can be appraised, the actual realized revenue from a sale will be uncertain. This trait makes fine art auctions a good field to apply and test empirically the findings about price risk. Conveniently, auction houses release price estimates prior to every auction. Sellers' price risk can then be defined as the degree to which realized revenue strays from the pre-sale appraisal by the auction house.

The news of a record-breaking bid to win a coveted Picasso a few years ago leads one to wonder what variables influence auction price risk. What can a seller do to control the amount of risk in a sale, and what effect will that control exert on the expected revenue? What types of bidders increase sellers' price risk and what kinds tend to have a stabilizing impact on price? How does the market in general affect price risk? And does the selection of one auction house over another determine the level of risk at the sale? Standard auction theory for the past four decades has been focused on format, buyers' risk, and bidding strategy and their impacts on total revenue, to the relative neglect of sellers' price risk in single object sales.

In Section 2, I explain the procedure of a typical fine art sale. Section 3 reviews important results of auction theory and empirical sale result analysis that are then used in Section 4 to derive a theoretical framework of auction price risk.

I have applied the results of that framework to the results of sales of Australian Aboriginal art over the last 10 years. I introduce the dataset and the methodology I used in compiling it in Section 5. In Section 6 I demonstrate how the auction price risk framework can be applied to improve our understanding of the types of bidders participating in a sale, historical changes in the marketplace, and competition amongst auction houses. I conclude in Section 7, where I offer my views on how the risk framework can be further generalized and extended to additional auction types.

2 Background on Fine Art Auction Sales

The fine art market is divided into two overlapping markets. The manufacturing side of the market consists of artists who produce art and sell it to dealers or collectors. The majority of the transactions in this market are fixed-price sales, not auctions. However, once art enters a private collection, it enters the second side of the art market, the resale sale market. There collectors and dealers wishing to sell an art work often engage an auction house to appraise the work, feature in in a pre-sale catalog, and then to execute a sale at an auction. If no bids are received, the lot is unsold and is returned to the seller. The price is determined solely by the demand at the auction, and fixed price sales are rare.

The fine art market features a set of products as diverse as any, and is usually subdivided into sub-markets by genre. Most auctioneers hold dedicated sales for each genre at regular intervals during the year, as determined by the amount of supply in the market. Those sales give the opportunity for people (bidders, potential sellers, connoisseurs, etc) interested in that genre to come together and observe the trends in the market and to participate in the sale.

Once an auction house selects the lots it is interested in featuring in the next sale, provenance is constructed. In-house art experts use their previous experience, art theoretical knowledge, and information about imminent demand to establish an estimate range (a low and a high price point) for each lot in the sale.

The estimate range is shown to the seller, who is given the opportunity to withdraw the lot from the sale if he disagrees with the estimate. Auction houses also allow the use of reserve prices (reserves), or a price that bidding has to exceed before the seller is legally bound by contract to sell the lot to the highest bidder.

Although having reserves is optional, auction theory has shown that there is an optimal reserve price for any auction. If no reserve price is set, the seller runs the risk of having to sell his/her possession at a price below his own personal valuation, and thus suffering economic loss. Furthermore, even when bidding exceeds the seller's personal reservation level, an optimal reserve price can maximize expected revenue from the sale. The majority of lots are sold with reserves, with the notable exception of estate or government sales, where liquidity considerations trump the risk of selling a lot below its market value.

While reserves are kept secret and auctioneers often engage in a bit of theater to prevent bidders from guessing it, each auction house has certain rules of thumb that give bidders an idea of the range for the reserve price. A common rule, for example is that the reserve price normally will not exceed some fraction of the low estimate and that lots with a low estimate under $1,000 will have no reserves at all.

A picture and a description of each lot, the information uncovered concerning provenance, and the estimate range are published in a pre-sale catalog that is distributed to potential bidders.

Although the pre-sale appraisal is not a particularly old tradition, I will argue that it is a revenue-maximizing part of the process. I will discuss theoretical and empirical evidence for this claim in the course of this thesis. The publication of biographical and authenticity information of each lot in a very affordable sale catalog has an obvious positive effect on the sale price by lowering the information costs that collectors would otherwise incur in their individual research of art works.[1]

The auction format used universally in fine art auctions is a modified version of the English ascending-bid auction. Like the English auction, bidders place bids as the price level is raised in steps by the auctioneer. However, due to the spread of bidding by phone or over the internet, bidders are not always aware of who exactly is still bidding and cannot observe at what price level another bidder dropped out of the auction. In this way, fine art auctions have become more similar to the Japanese ascending auction, where bidders can see how many (but not which) competitors are still participating in the bidding. This change is important as it means that bidders cannot observe when a famous collector or a museum stops bidding and will therefore feel a weaker Winner's Curse than at a fully public auction.

Auction houses schedule sales based on two considerations: the amount of available supply offered for sale and the amount of demand for the genre. The success of a sale is highly variable – some sales raise much more than expected while other do not. Auctioneers measure demand through the buy-in rate, the percentage of lots that fails to sell because no bids cleared the reserve price.

[1] The convenience value of the catalog lowers the barriers to entry into the auction, increasing the number of bidders, and raising everyone's valuations by removing uncertainty and information costs. This also raises the revenue a seller can expect.

3 Literature Review

3.1 Auction theory

Much of early auction theory work relies on the assumptions of the private values model. Because of its intuitive simplicity, I too will begin my analysis of risk with the private values model. That framework involves a single lot offered to n bidders, each of whom knows their own valuation of the lot, but is not aware of how the others are valuing the same work. Each bidder is assumed to draw their valuation from a continuous distribution with the cumulative density $F(v)$ that is known to all participants. Because each draw is Independently and Identically Distributed (IID) relative to other draws, the valuation of one bidder has no influence on the valuation of another. The function $F(v)$ is assumed to have a domain of $[v_*, \bar{v}]$, and to be continuously differentiable in that domain with a first derivative $F'(v)$. Because $F(v)$ is the cumulative distribution function of bidder valuations, $F(\bar{v}) = 1$ and $F(v_*) = 0$; that is, I assume that bidders whose valuations lie below the cut-off point v_* have no chance of actually winning the auction. People with valuations below v_* will find that even if they manage to "steal" the item at their low valuation, they will gain too little from the tiny chance they have of actually taking the lot home. To verify that assumption, at most auctions that I have attended, even when there is no reserve price, the auctioneer almost never starts the bidding at \$0.00 but at some price higher. On the other hand, there is a price level $\bar{v}$ that guarantees that bidder the lot. While it is technically an unrealistic assumption, it reduces complexity in many proofs and allows the use of algebraically simple density functions. The seller is also assumed to have a private value v_0 for the lot that is strictly bound by the same domain as the bidders' density functions.

Riley and Samuelson elegantly showed in 1981 that the expected revenue to the seller from an auction sale in the private values framework is

$$n\int_{v_*}^{\bar{v}}\Big[vF'(v)-1+F(v)\Big]F(v)^{n-1}dv \tag{1}$$

Because the above equation applies to any auction that abides by the two above mentioned assumptions, it is a part of the proof of the Revenue Equivalence of auctions, a theory that shows that all types of auctions in the private values model lead to the same expected revenue to the seller [Riley and Samuelson, 1981]. I will use that result in extending the expected revenue result to art auctions.

The concept of seller's auction risk is usually discussed in the framework of auction type. Even though revenue equivalence guarantees equality of expected revenue, some auction types are less risky than others. For instance, the second-price English auction is considerably less risky to the seller than the first-price descending or sealed-bid auction [Krishna, 2002]. The seller should always choose the second-price auction, which explains why second-price auctions are by far the most popular format for indivisible single-unit auctions.

The exact distribution $F'(v)$ is usually left in general form. Because it has a limited domain, it cannot be the normal distribution or any other infinite domain function. The case of the uniform distribution was used by Riley and Samuelson to derive some basic conclusions in their paper on optimal auctions as well as in many basic applications of auction theory [Riley and Samuelson, 1981]. To derive a framework within which to be able to discuss sellers' price risk, I will use two types of distributions, $F(v)=\frac{v^a-v_*^a}{\bar{v}^a-v_*^a}$ and $F'(v)=(v-v_*)(\bar{v}-v)$. The first class includes the uniform case as well as increasing and decreasing functions. The second case rests on the intuition that valuations towards the middle of the range are more likely to occur than the extremes.

In the revenue maximizing case, it has been shown that the seller should set reserves at a level strictly higher than v_0 [Klemperer, 1999]. This is derived from the above equation for expected revenue from the sale. Because of the above-mentioned rules of thumb for reserves, the fine art market seems in the face of it to be inefficient, as the seller's reservation value could be quite a bit higher than the low estimate. This in turn can result in unwarranted costs if it increases the buy-in rate. A further investigation in Section 5 shows how other practices by auctioneers compensate for the restrictiveness of the rules of thumb.

Does the fine art market fit the private values assumptions? While each bidder appears to assign a private value component to his/her valuation – whether the painting matches their kitchen or maybe reminds them of childhood will have a large influence on their willingness to pay for it – bidders are clearly also influenced by the outside world. Literature specific to the art market has attempted to enumerate and quantify a number of those influences.

3.2 Art market

Fine art auctions typically see two distinct classes of bidders: dealers and private collectors. Dealers are the more gain-oriented, and are looking for lots whose resale value they believe could be higher at a later date. They are not likely to win a bidding war against a private collector who has his sights set on a piece [Singer and Lynch, 1994].

Private collectors are not only influenced by their expectation of the resale value, but also by their personal private value for the lot. That value is sometimes referred to as the "use" value because it corresponds to the utility that the owner receives from ownership of art before any investment returns. A collector is more likely than a dealer to bid significantly higher than the expected price for the lot because he/she may have a high use value. Collector tastes are capricious, and hard to measure, adding more uncertainty to the valuation. In application to auction price risk, I expect that lots bought by collectors should exhibit more risk than lots bought by dealers, and for that reason a seller should try to discover the composition of the likely bidders before making the decision to hold an auction.

Fine Art has always been a luxury good in the Western world, and used as an investment vehicle by many. As the world economy grows, the demand for masterpieces outpaces their supply (and for some genres, like Old Masters, the world stock is actually shrinking due to depreciation and loss to fire and thieves). The degree to what buyers are interested in works as investment assets versus as art works to be enjoyed for their æsthetic appeal depends on the recent historic trends in art valuations versus their closes financial benchmarks.

Art literature has made copious use of Rosen's formulation of the hedonic price index [Rosen, 1974]. Unlike financial assets, artworks are resold very infrequently, making it difficult to determine the growth in valuations on year-on-year basis. By regressing sale prices by the various properties of art, scholars have proxies for value growth rates.

The most recent study of the largest cross-section of the art markets has been by Mei and Moses, who showed that while art performed well as a component of a diversified portfolio, its risk-adjusted returns underperformed its benchmarks [Mei and Moses, 2002]. Mei and Moses showed that owning art as an investment subjects the holder to greater relative risks to returns compared to other financial instruments. When combined with the formulation of auction risk, the full picture of risk that accompanies fine art is exposed.

A hedonic regression study of the Australian Aboriginal art market appeared in 2005. The analysis, by Higgs and Worthington showed that since the mid 1990's prices of Aboriginal art have grown at a high rate after the slump in the early 1990's [Higgs and Worthington, 2005]. The breakdown of their (albeit much larger) dataset by artist and price is consistent with the data in my own dataset.

4 Theoretical Framework

4.1 Derivation of Variance in General Form

The expected payment of any bidder of type v_1 at a private value auction is defined as

$$P(v_1) = v_1 F(v_1)^{n-1} - \int_{v_*}^{\bar{v}} F(x)^{n-1} dx \tag{2}$$

The expected revenue from bidder 1 ($\bar{p_1}$) for the seller is then just the sum of the payments times the bidder's probability of winning at that payment,

$$\bar{p_1} = \int_{v_*}^{\bar{v}} P(v) F'(v)^{n-1} dv \tag{3}$$

substituting in the $P(v)$ and switching the order of integration, I have

$$\bar{p_1} = \int_{v_*}^{\bar{v}} \Big[vF'(v) - 1 + F(v) \Big] F(v)^{n-1} dv \tag{4}$$

Equation 4 is the first moment of revenue from a single bidder. To compute the the variance of an individual bidder's revenue, I can use the second moment,

$$\sigma_1^2 = \int_{v_*}^{\bar{v}} \left(v_1 F(v_1)^{n-1} - \int_{v_*}^{v_1} F(x)^{n-1} dx \right)^2 * \\ F'(v_1) dv_1 - \bar{p}^2 \tag{5}$$

And because all bidders are symmetric and independent of one another, the variance of seller's revenue is simply the sum of the variance of each bidder, or $n\sigma_i^2$. Because variance is not expressed in applicable units, I will mostly refer to the standard deviation of the seller's revenue, which is just the square root of the variance.

A challenge to a comprehensive framework of auction risk is the double integral in the solution for variance. It makes calculating variance for complex density functions impossible. However, I can evaluate variance of expected sale revenue by using algebraically simple distributions for $F(v)$. Studying the results from these distributions is still useful as they may show patterns that match these we observe in real-life auctions. How differences in density functions create differences in variances is important knowledge for an auctioneer trying to structure a transaction that fosters a particular form of auction price risk.

4.2 Distributions

4.2.1 Polynomial-type density functions

The first type of density distribution studied is the simple polynomial defined over domain v_* to $\bar{v}$[2] of form

$$F(v) = \frac{v^a - v_*^a}{\bar{v}^a - v_*^a} \tag{6}$$

where a is a real variable determining the exact type of distribution used. When $a = 1$, the resulting density is the uniform distribution. For $a > 1$, the density is an increasing polynomial with a positive second derivative. These densities suggest that most common valuation will actually higher than the average valuation, *e.g.* that the

[2]For the sake of simplicity, I will omit the full piecewise definitions of distributions, but the reader should keep in mind that functions $F(v)$ and $F'(v)$ are equal to 0 when $v > \bar{v}$ or $v < v_*$

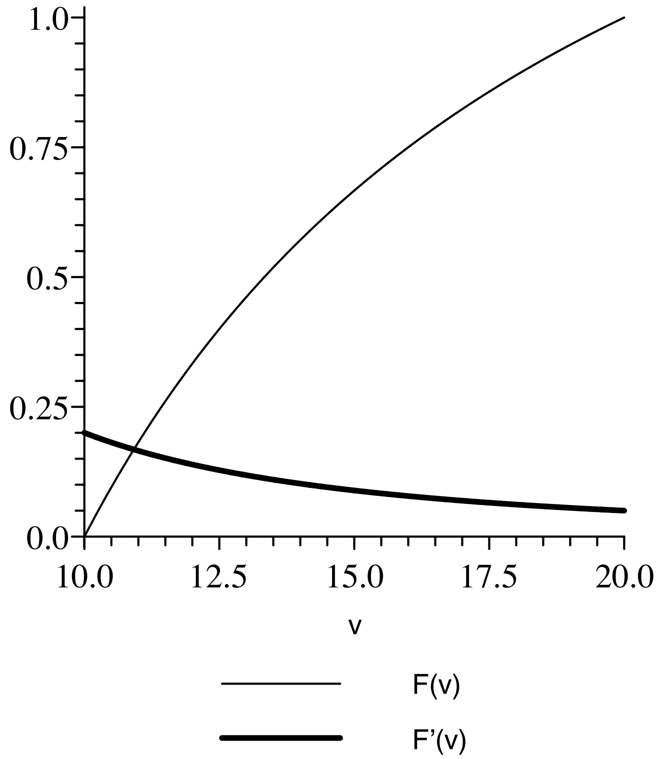

Figure 1: Density functions when $a = -1$, domain [\$10 . . . \$20]

bulk of bidders will have high valuations but a few will have much lower valuations.

The density $a = 0$ is not defined, but densities the correspond to $a < 0$ are decreasing rational functions with positive second derivatives. These densities suggest the opposite of the last case, that most bidders have low valuations but a few high v bidders make the average higher (Figure 1).

4.2.2 Negative quadratic distribution

An alternative distribution that is simple enough to be solvable in general form is the negative quadratic distribution with the probability density of

$$F'(v) = \frac{-6(v - v_*)(\bar{v} - v)}{-\bar{v}^3 + 3v_*\bar{v}^2 + v_*^3 - 3v_*^2\bar{v}} \tag{7}$$

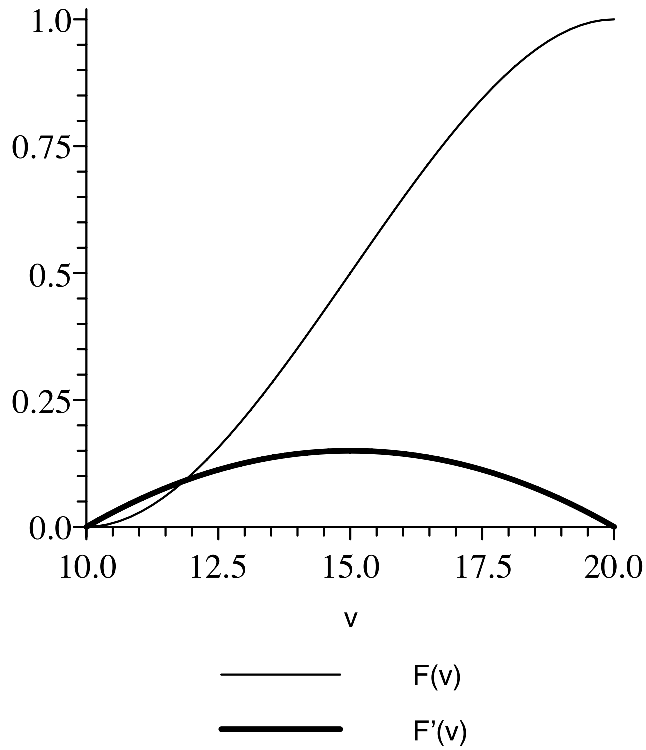

Figure 2: Density functions for quadratic distribution, domain [\$10 . . . \$20]

The quadratic density distribution in Equation 7 is a more intuitive distribution than the polynomial distribution. As a negative quadratic function, it peaks in the middle of its domain, as opposed to the beginning or the end. In this way, it is similar to the normal distribution, the most widely occurring distribution in observed phenomena. It makes sense that while some people may have valuations far outside the average due to their individual tastes, the most common opinion should not lie at one of the extremes, but somewhere in the middle (Figure 2). Results from this distributions turn out very similar to these of the $a > 0$ distributions, but I include it in my analysis for completeness.

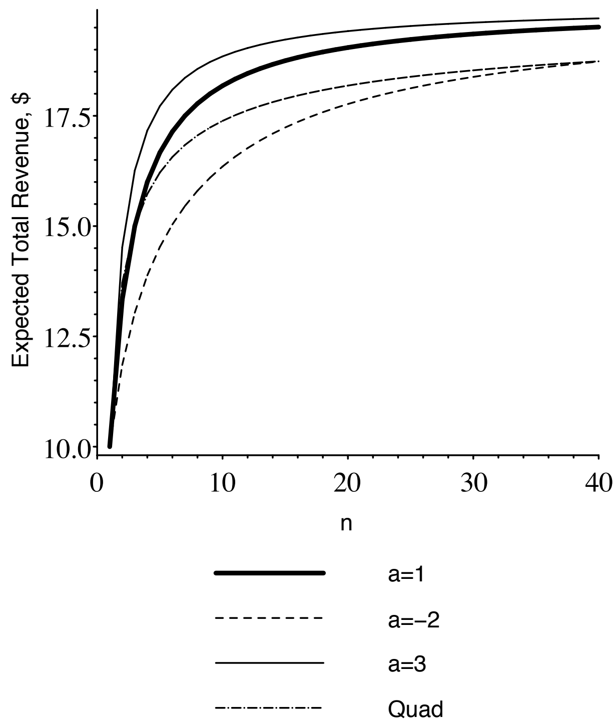

Figure 3: Expected Revenue by number of bidders n

4.3 Impact of auction variables on expected revenue

By substituting density functions into Equations 4 and 5, I can plot the behavior of total expected revenue and price variance in response to changes in the independent variables of the number of bidders and the relative size of the density domain $\bar{v} - v_*$.

4.3.1 Number of bidders

Figure 4 plots the the relationship between the number of bidders and each bidder's expected payment. The first marginal bidders have a large positive effect on each existing bidder's expected payment, as competition amongst bidders reduces the opportunity for bidders to realize gain from the auction, that is to purchase a lot for a price significantly below their valuation. That effect persists even as the number of bidders increases, but is counteracted by the falling probability of actually winning the sale. The marginal change in expected payment approaches zero as number of bidders continues to increase.

The total expected revenue to the seller is just the sum of all bidders' expected payments. In a symmetric auction (like the fine art auction), all bidders have the same payment expectation and calculating total expected revenue is trivial. Figure 3 plots the expected revenue against the independent variable of the number of participating bidders.

The expected revenue is an increasing function of the number of bidders. Since an additional bidder can never send a negative signal to the other bidders, his/her entrance will increase revenue if he draws the highest valuation of the bidders who are already participating, or will have no effect on revenue if he draws a valuation lower than the highest bid. Expected revenue has a negative second derivative – the marginal effect of an extra bidder decreases and gradually approaches zero, just like each bidder's expected payment. The entrance of a new bidder adds one extra payment to the pool, but also depresses every single bidder's expected payment. In the limit of infinite number of bidders, the seller's expected revenue is $\bar{v}$.

4.4 Impact of auction variables on price risk

4.4.1 Number of bidders

Figure 4 plots the relationship of the standard deviation of bidder 1's payment and the number of participating bidders. Standard deviation peaks when the number of bidders reaches 5, and then begins to decline. The decline is much steeper for the $a \geq 1$ and quadratic densities than it is for the $a < 0$ densities, which remain virtually flat as new bidders enter the market. The overall level of variance is also significantly

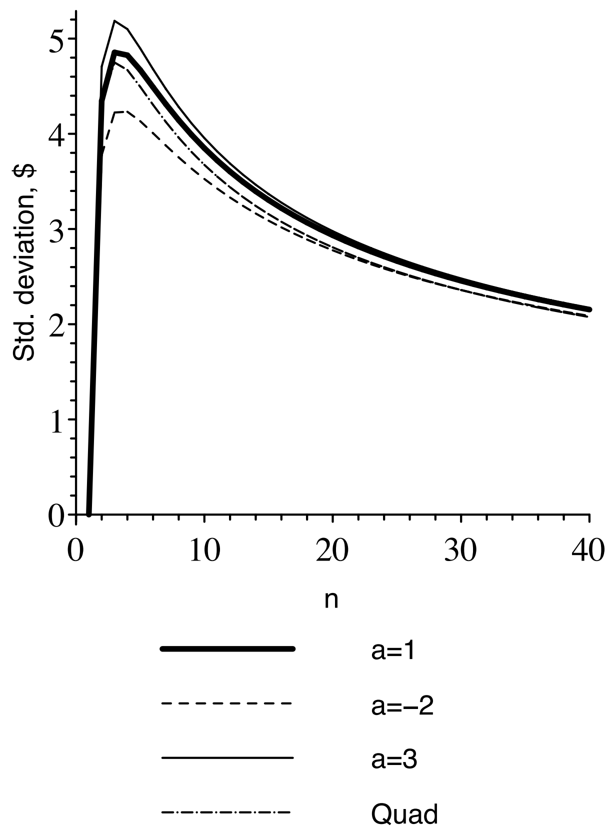

Figure 4: Standard deviation of bidder 1's payment by number of bidders n

lower.

While each bidder becomes more certain of his/her chances of winning as more bidders enter the auction, the seller's risk continues to grow. Figure 1 plots the relationship between the standard deviation of seller's revenue and the number of participating bidders. The $a < 0$ densities produce a much less risky environment for the seller than the $a > 0$ distributions.

Note that the standard deviation is 0 at points when there are no participating bidders and when there is an infinite number of bidders. However, it is never close to zero in between.

In the case that the seller wants to minimize risk but is expecting a poor turn out

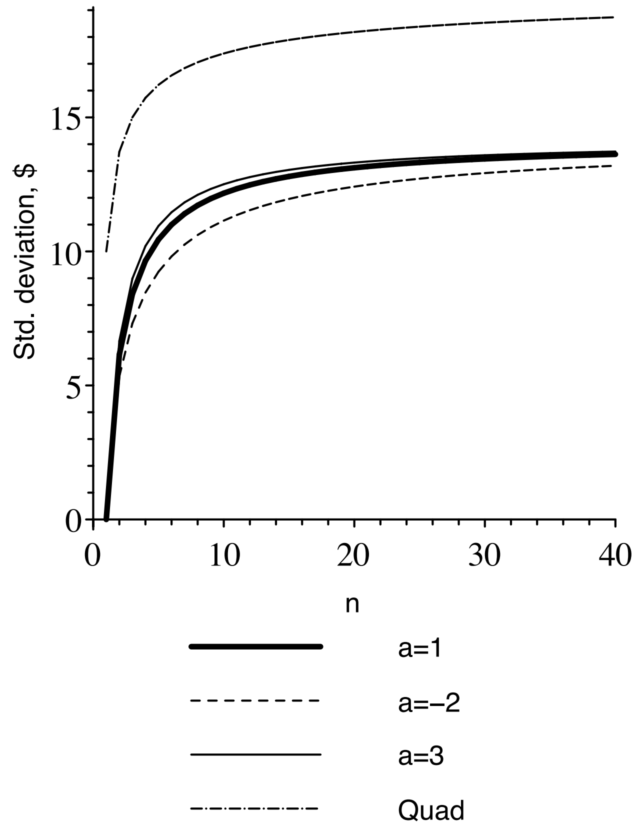

Figure 5: Standard deviation of revenue by number of bidders n

to the auction, it will be in his/her interest to cater to bidders who will draw from the $a < 0$ density. A marginal bidder from that distribution will not result in as large an increase in risk as a bidder from a $a > 0$ distribution, but will still have a favorable effect on total revenue.

4.4.2 Size of density domain

Auction price risk depends on the relative size of the density domain. Whether the highest possible valuation $\bar{v}$ is twice or three times the level of the lowest possible valuation v_* will affect the price variance of a seller who is selling to bidders from either of the two densities.

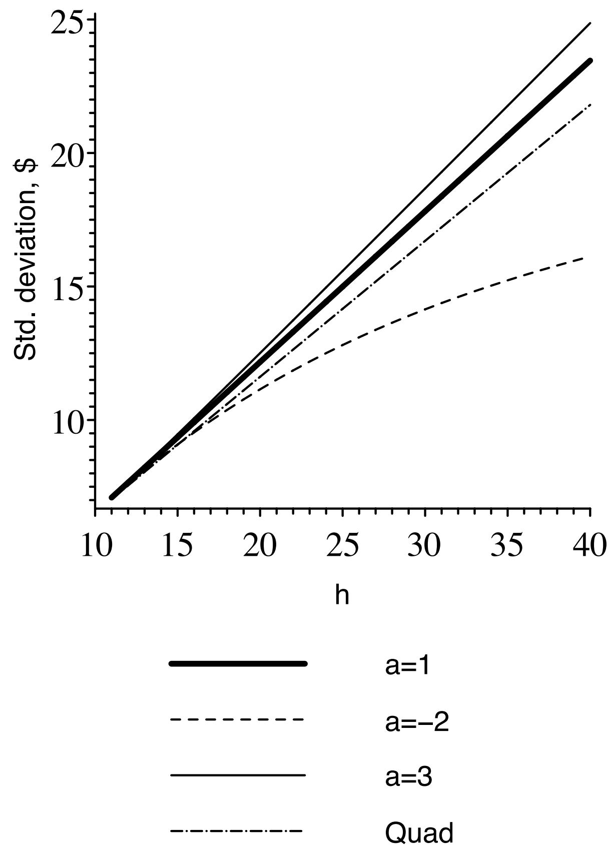

Figure 6: Standard deviation of revenue by size of $\bar{v} - v_*$

Figure 6 plots variance of several densities against different levels of $\bar{v}$. For all functions, v_* is assumed to remain constant at $1. That way, a change of $1 in $\bar{v}$ will represent a marginal increase in the size of density domain of 100%.

The $a = 1$ distribution shows a perfectly linear relationship with domain size. A 1% increase in $\bar{v}$ will spur a smaller, but fixed percent increase in seller risk. $a = -1$ density function, on the other hand, shows a declining effect of domain size on risk as domain size grows. The density function presupposes such a clustering of bidders in the low end of the valuation domain that a stretching of the domain does not result in equivalent number of new high valuation bidders to push up risk. At the point where

$\bar{v}$ is ten times the v_*, the risk involved in selling to bidders from the $a = 1$ density is twice that of selling to bidders from the $a = -1$ density, which is also seeing much smaller marginal growth in risk from growth in valuation domains.

5 Data

Fine art auctions make good subjects for analyzing seller's auction risk because of the existence of auction house estimates that capture the collective expectation of auction results. In choosing an art market for the empirical section of this study, I looked for a relatively small but growing market that will have bidders who will assign the most value to the private value component of their valuation. Art markets that assign a large part of the valuation equation to provenance would not have as good a fit to the private values model as markets where the important painters are still alive and authenticity is rarely questioned. The other major concern in the selection process was the number and frequency of resales. A large number of resales would indicate that dealers, speculators, and investors have come to dominate the market, the exact type of bidders who will focus more on the expectation of the resale value rather than the private "use value" of art.

That criteria led me to the Australian Aboriginal art market, a relatively new market that first appeared in the early 1970's, and really took off in the 1990's. While sales are usually held once or twice a year in either Sydney or Melbourne, all the major auction houses established Aboriginal art as either an independent department or combined with Oceanic art to form a department. Auction catalogs regularly list estimates in Australian dollars, U.S. dollars, and Euro in an attempt to cater to the international clientele. Viewed by many as a growth market, the medium term strategic importance of Aboriginal art to auction houses is the challenge to keep sales efficient and fair to both buyers and sellers.

Because the market is relatively young and has few masterpieces well-known in the general population, I hypothesized that the number of resales would be few and they would be far between. I was later proven correct in that assumption; of the 5,000 sales in my dataset, there were fewer than a hundred sales that matched in both title and artist. A thorough inspection of these hundred would likely find that even fewer were actual resales. While a greater percentage of sales would qualify as resales if the data included more of the early sales, sellers who held the lot for at least five years are more likely to be connoisseurs or collectors rather than investors, dealers or speculators who are buying art for its expected resale value rather than its æsthetics.

The dataset consists of transcriptions of pre-sale auction catalogs from 1995 to March 2006, combined with the realized results from said auctions. Only Australian Aboriginal art is included, and only when it was sold at an auction sale dedicated to the genre. The rationale for these restrictions is to make sure that we capture the behavior of a set of bidders in a particular market. Since I am partially making use of the assumption that all bidders are drawing from a common distribution of valuations, limiting collection to events as similar as possible to each other is essential to maintaining the viability of that assumption.[3]

The data shows a vibrant marketplace with competition among the six participating auction houses.[4] For the purposes of this study, I will assume that sales by all three can be treated identically.. While Sotheby's is consistently the leader by market share, Menzie's and Christie's vie for the second place. In recent years Menzie's has managed to outsell Christie's, which in March 2006 announced that it will be exiting

[3] I have focused on traditional Aboriginal paintings, which dominate modern (mostly urban) paintings and other art, but I have included in my database painted statues and weapons, and supports of canvas, board, and bark.

[4] In this dataset, *Menzie's* actually stands for three different auction house brands that are owned by the same corporate entity: *Lawson∼Menzie's*, *Deutscher∼Menzie's*, and *Lawson's*.

the Australian Aboriginal art market.

The league table (Figure 7) by artist is lead by Rover Thomas and Emily Kame Kngwarreya, the two artists who made the genre famous by placing their works into prestigious museums of modern art. While various younger artists are beginning to emerge as buyers search for a new generation of artists, the rest of the marketplace is split among a relatively small number of other established artists. Sotheby's controls the lion's share of sales by the top 20 artists in the table, signifying that it is the *de facto* auction house of choice for owners of best-selling Australian Aboriginal art. Figure 8 breaks down sales by the year and the auction house. I will investigate below what factors other than brand could be making Sotheby's attractive to sellers of high worth Aboriginal art.

6 Empirical Specification

Auction price risk is measured as the standard deviation of the realized revenue to the seller and the expectation of that revenue given the information the seller had about the shape and domain of the distribution $F(v)$ and the expected number of bidders who would participate.

However, we cannot directly observe the expected revenue to the seller, as we do not know the common density distribution, its domain, or the expected number of bidders. Could the high estimate perhaps serve as a proxy for this expectation?

The log of the high estimate has a very strong correlation with realized revenue. In percentage terms, winning bids tend to stray farther from the high estimate for more expensive lots, but winning bids cluster closely around the 45 degree line in a log-log construct (Figure 9).

In the dataset, the low estimate emerges as being on average two-thirds of the high estimate. While it is a good predictor of the final bid, the high estimate requires no coefficient and no regression to have a good fit with the data. It is peculiar, however, that auction houses call their "real" estimate the high estimate, instead of targeting the mean of the high and the low estimates. The explanation is that the buyer's premium, a $10-20\%$ surcharge on the winning bid that is collected by the auction house as a fee for organizing the sale.

Since I have shown that the high estimate can serve as a very strong proxy for the expected revenue from a sale, I can use it to construct a proxy for the observed sale price risk. The square of the difference between the high estimate and the realized price (with buyer's premium) is a proxy variable for the observed standard deviation in the sale.

Figure 10 charts the distribution of observed deviations in all sales as they compare to the high estimate. Deviations tend to cluster around a median value and have a log-normal distribution that the logarithm converted to a normal distribution. Figure 11 shows a plot of how the average and median of the deviation function changed over the years. The mean remained remarkably stable over the years, suggesting that even as valuations have risen, the underlying driver of variance, the density function $F(v)$, did not experience large changes.

Another way to analyze price variance would be by auction house. As mentioned earlier, Sotheby's has a leading market share in this market, followed by Menzie's and Christie's.

To measure the effect an auction house has on the variance of sales it executes, I ran the following regression:

$$\log \sqrt{(X-\mu)^2} = \Big(\beta_0 + \sum_{k=1}^{m-1} \beta_k\Big) \log \mu + \beta_{m+1} \tag{8}$$

where $\beta_1..\beta_{m-1}$ are Boolean variables that identify the auction house that executed that particular sale, X is the realized sale price (only lots that sold were included), and μ is the published high estimate. m is the number of auction houses in the dataset, but one is omitted from the regression (in this case Sotheby's) to prevent perfect multicollinearity. To determine the level of price risk associated with each auction house one examines the statistical significance of its affiliated β_k coefficient. A significantly positive or negative coefficient signifies that the auction house has a risk level that is significantly different from the base level suggested by the omitted auction house. The coefficient $\beta_0 + \beta_k$ calculates the exact relationship between the high estimate and the deviation for that auction house. The results of the regression are tabulated in Figure 12.

The regression reveals two concepts. The first one is illustrated in Figure 13, where a clearly linear pattern emerges between a percentage change in the high estimate and a corresponding change in the deviation. The second revelation is that the only auction houses statistically different from Sotheby's level of risk were Christie's and Shapiro, both of whom had lower price risk than Sotheby's. This suggests that if you expect the number of bidders, and their valuation densities to be the same everywhere, and if all houses estimate the lot at the same price, then the seller will be best off selling at Christie's or Shapiro as they would minimize the seller's price risk. The choice of the auction house, however, is much less important to seller's risk than other factors.

The statistically significant differences among auction houses are reflections of the different number and types of bidders that each one attracts. Christie's had the largest mean high estimate selling price of $\$16,500$, over $\$6,000$ more than Sotheby's average, and therefore could have attracted fewer bidders than its competitors.

Because all the coefficients in Figure 12 were less than one, the regression suggests that either the density function, the number of bidders, or the relative size of the domain of the density function changes as the high estimate increases. If both variables were to stay constant, then the growth in deviation would correspond one-to-one to growth in the high estimate and in v_*,$\bar{v}$. The decline in relative size of the deviations can be attributed to either a decline in the number of bidders that accompanies growth in estimates or to compression in the relative size of the domain $[v_*, \bar{v}]$ that occurs continuously with growth in $\bar{v}$.

Both explanations are plausible. One can imagine an auction sale taking place in a large ballroom. The people sitting upfront wearing designer suits and talking about their newest Ferrari can most definitely afford to bid for every lot that is offered on sale. However, the people hanging back by the buffet table are probably only looking a fraction of the lots offered. Therefore, while everyone in the room is a potential bidder for a cheaper work, only a few participate in the sale of a masterpiece. The second explanation is less obvious, but just as intuitive. Higher-worth art works have typically been exhibited and studied more in comparison to cheaper works. That means that their standing and æsthetic worth among the other works by that artist has been determined to a greater extent than for a lesser-known work. While the lesser-known work may one day be "discovered" and become much more valuable, a well-known work will carry less movement in valuation. The greater certainty in estimating the lot's æsthetic value allows auction houses to make better estimates and has a compressing effect on the density domain. Bidders whose valuations laid closer to v_* will feel that the available information justifies paying a little extra, while bidders whose valuations laid closer to $\bar{v}$ believe that the upside risk is not as high as they may previously hoped (for example, the work has already been considered and passed over at a major exhibition), and would lower their valuations accordingly.

To test the second hypothesis, I created a Boolean variable α that was 1 when the lot was produced by one of the top 20 best selling artists (the same ones listed in

Figure 7) and ran the following regression:

$$\log \sqrt{(X-\mu)^2} = (\beta_0 + \beta_1 \alpha) \log \mu + \beta_2 \tag{9}$$

The results disproved the hypothesis. The regression found β_1 to be 0.01, meaning that works produced by best-selling artists are associated with more auction risk than works produced by lesser known artists.

7 Conclusion

Using the fundamental relationships of auction theory, I derived an equation governing the variance of seller's total revenue. The properties of revenue variance were analyzed inductively by plugging in several density distributions into the variance equation, and by analyzing their first order partial derivatives in respect to the number of bidders and the relative size of the domain of the density function. Assuming that there exists a set of classes of bidders, each drawing valuations from a particular density function (an example of such classs might be the private collectors, the art investors, and the art dealers), then the composition of bidders by these classes will play a pivotal role in determining price variance.

The application of the general formulation of revenue variance revealed a fundamental trade-off between risk minimization and revenue maximization. A marginal bidder has a positive effect on both the expected revenue and the standard deviation of revenue, meaning that it is impossible to optimize both at the same time.

The empirical application of revenue risk to data from sales of Australian Aboriginal art revealed differences in observed variances of art sold in different years, in different auction houses, and in art from top-selling Aboriginal artists and that from lesser-known artists.

This thesis can be extended further by examining additional distributions.

References

[Higgs and Worthington, 2005] Higgs, H. and Worthington, A. (2005). Financial returns and price determinants in the australian art market, 1973-2003*. *Economic Record*, 81(253):113–123.

[Klemperer, 1999] Klemperer, P. (1999). Auction theory: A guide to the literature. Technical Report 9903002, Economics Working Paper Archive at WUSTL. available at http://ideas.repec.org/p/wpa/wuwpmi/9903002.html.

[Krishna, 2002] Krishna, V. (2002). *Auction Theory*. San Diego : Academic Press.

[Mei and Moses, 2002] Mei, J. and Moses, M. (2002). Art as an investment and the underperformance of masterpieces. *American Economic Review*, 92(5):1656–1668.

[Riley and Samuelson, 1981] Riley, J. G. and Samuelson, W. F. (1981). Optimal auctions. *The American Economic Review*, 71(3):381–392.

[Rosen, 1974] Rosen, S. (1974). Hedonic prices and implicit markets: Product differentiation in pure competition. *The Journal of Political Economy*, 82(1):34–55.

[Singer and Lynch, 1994] Singer, L. and Lynch, G. (1994). Public choice in the tertiary art market. *Journal of Cultural Economics*, 18(3):199–216.

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

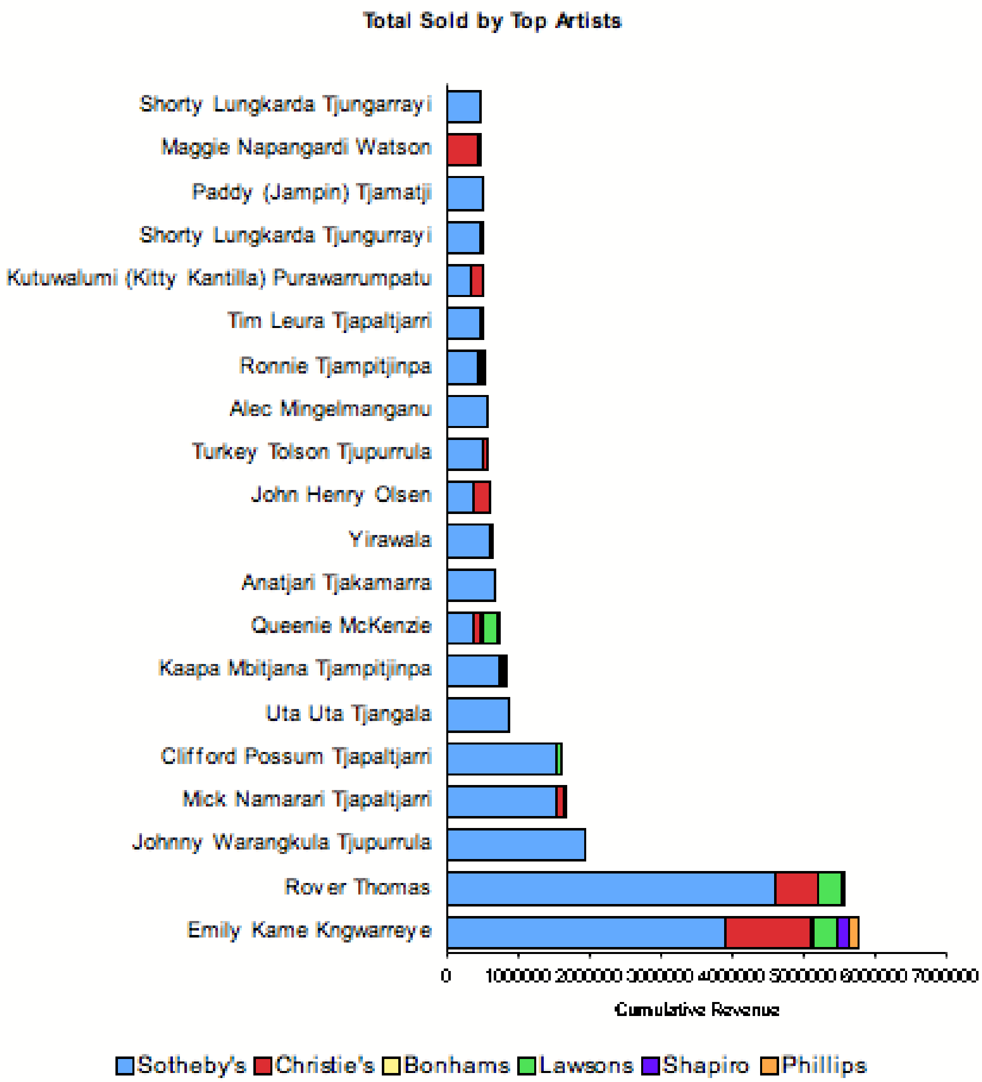

Figure 7: League table of top-selling artists

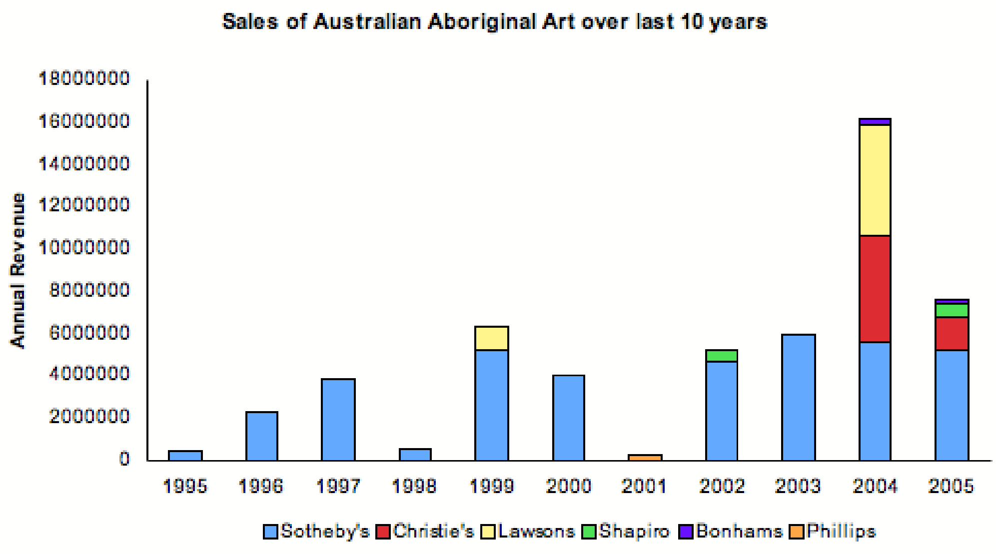

Figure 8: League table of highest volume auction houses

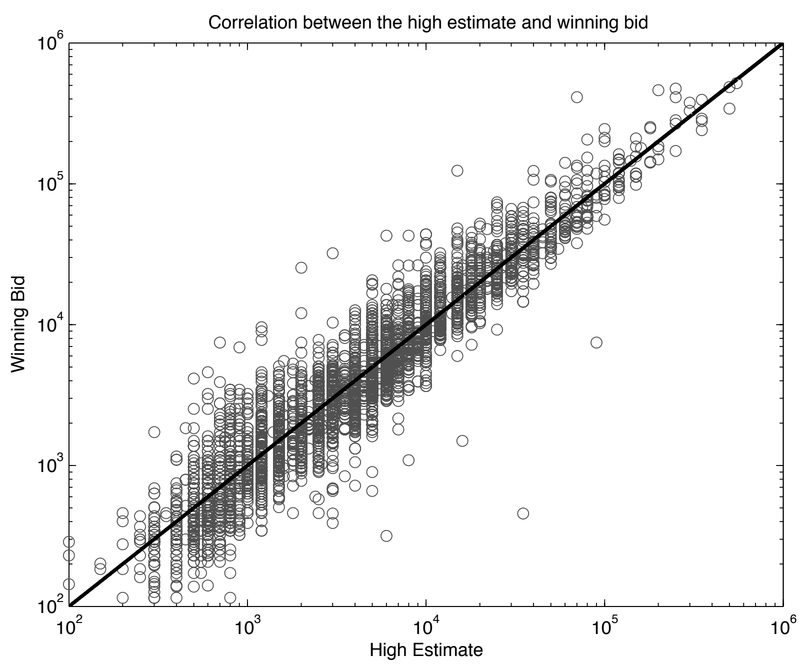

Figure 9: Log-log relationship between μ and X

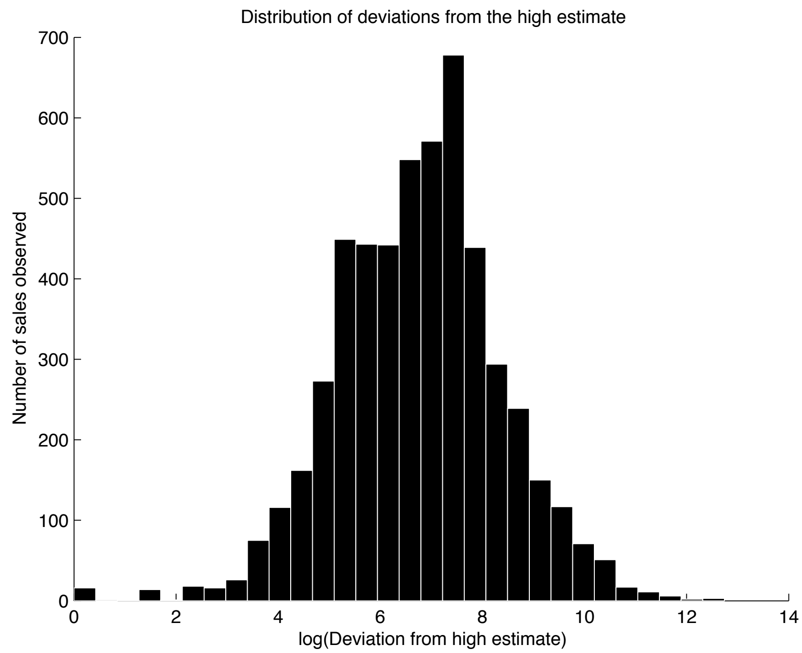

Figure 10: Histogram of logs of the observed deviations

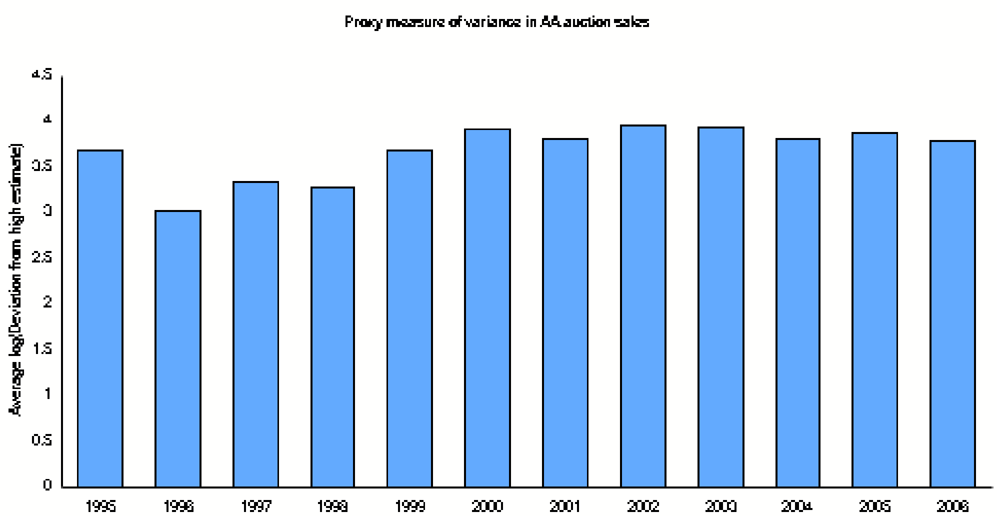

Figure 11: Historic trends in observed standard deviations from μ

Rank	Auction House	$\beta_0 + \beta_k$
1	Phillips	0.94
2	Sotheby's	0.93
3	Lawson's	0.92
4	Shapiro	0.91
5	Bonham's	0.91
6	Christie's	0.90

Figure 12: Ranking auction houses by growth rate of observed deviations

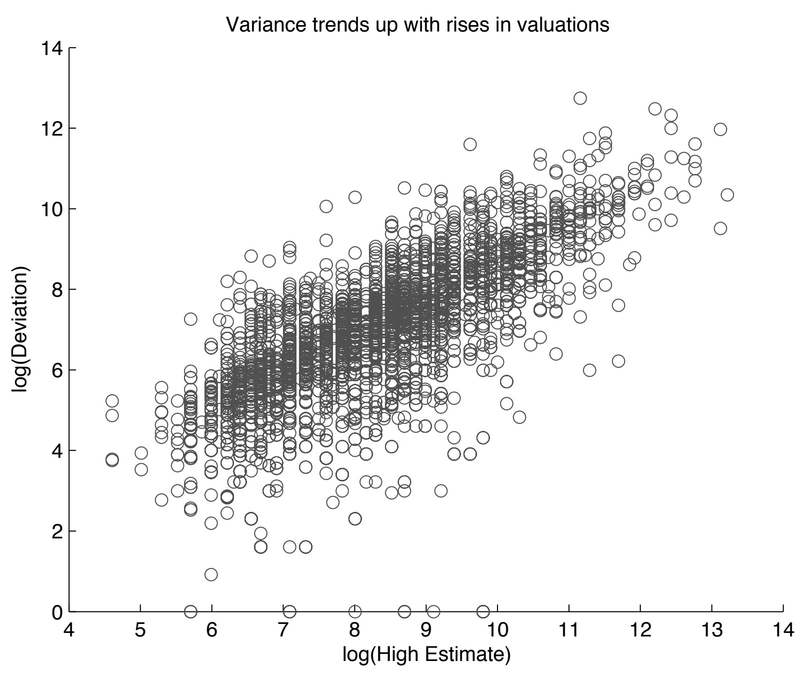

Figure 13: Log-log relationship between μ and $\hat{\sigma}$

Appendix A: List of auction sales included in data set

Sum of Fmt Sale Price AUD		
Fmt Auction Year	Sale Of	Total
1995	Sotheby's Melbourne: Sunday, June 18, 1995	$226,149
	Sotheby's Melbourne: Tuesday, November 28, 1995	$250,815
1995 Total		$476,964
1996	Lawsons: Tuesday, May 21, 1996	$1,650
	Sotheby's Melbourne: Monday, June 17, 1996	$1,407,221
	Sotheby's Sydney: October 27-28, 1996	$878,448
1996 Total		$2,287,319
1997	Lawsons: Tuesday, December 9, 1997	$9,702
	Sotheby's Melbourne: June 30, 1997	$2,727,663
	Sotheby's Sydney: November 9, 1997	$1,083,740
1997 Total		$3,821,105
1998	Sotheby's Sydney: November 9, 1998	$594,758
1998 Total		$594,758
1999	Lawsons (Deutscher-Menzies): June 29, 1999	$1,109,690
	Sotheby's Melbourne: June 28, 1999	$2,900,559
	Sotheby's Melbourne: Monday, November 22, 1999	$2,157,188
	Sotheby's Melbourne: Tuesday, November 23, 1999	$195,270
1999 Total		$6,362,707
2000	Sotheby's Melbourne: Monday, June 26, 2000	$3,939,384
	Sotheby's Melbourne: Tuesday, June 27, 2000	$76,303
2000 Total		$4,015,687
2001	Lawsons: Monday, May 21, 2001	$2,450
	Phillips Sydney: Monday, July 30, 2001	$301,990
2001 Total		$304,440
2002	Shapiro Auctioneers: Thursday, May 9, 2002	$237,930
	Shapiro Auctioneers: Tuesday, December 3, 2002	$237,867
	Shapiro Auctioneers: Wednesday, May 8, 2002	$63,215
	Sotheby's Melbourne: Monday, June 24, 2002	$4,707,100
2002 Total		$5,246,112
2003	Sotheby's Sydney: Monday, July 28, 2003	$4,748,768
	Sotheby's Sydney: Tuesday, July 29, 2003	$1,180,336
2003 Total		$5,929,104
2004	Bonhams & Goodman: Wednesday, November 17, 2004	$225,785
	Christie's Sydney: Tuesday, November 30, 2004	$2,027,559
	Christie's Sydney: Tuesday, October 12, 2004	$3,009,240
	Lawsons (Lawson-Menzies) Sydney: May 25, 2004	$1,535,424
	Lawsons (Lawson-Menzies): May, 2004	$1,542,135
	Lawsons (Lawson-Menzies): November, 2004	$2,244,330
	Sotheby's Melbourne: Monday, July 26, 2004	$5,013,978
	Sotheby's Melbourne: Tuesday, July 27, 2004	$560,740
2004 Total		$16,159,191
2005	Bonhams : August 30, 2005	$102,997
	Bonhams : November 23, 2005	$98,449
	Christie's Melbourne: Tuesday, August 30, 2005	$1,545,718
	Shapiro Auctioneers: Sunday, March 20, 2005	$375,066
	Shapiro Auctioneers: Wednesday, October 26, 2005	$267,740
	Sotheby's Melbourne: Monday, July 25, 2005	$4,336,285
	Sotheby's Melbourne: Tuesday, November 15, 2005	$774,360
	Sotheby's Melbourne: Wednesday, November 16, 2005	$125,160
2005 Total		$7,625,775
2006	Bonhams : March 29, 2006	$108,269
2006 Total		$108,269
Grand Total		$52,931,430