The Informational Efficiency of the Art Market

William N. Goetzmann

Yale School of Management

Forthcoming, *Managerial Finance*

Abstract

Two measures of informational efficiency are applied to the market for paintings. The first is a measure of market efficiency as captured by serial dependency in returns. The serial correlation in an index of art returns suggests the possibility of persistent trends in the art market, however there is no empirical evidence that these trends can be easily exploited. The second is a measure of "price risk," or instantaneous uncertainty about the immediate resale value of a work of art. The magnitude of the price risk suggests that there is a major role for dealers in the art market. Using historical data, I find that the price risk has been declining since the beginning of the painting market, indicating increasing informational efficiency.

Please address correspondence to:

William N. Goetzmann
Yale School of Management
Box 208200
New Haven, CT 06520-8200
will@viking.som.yale.edu

I wish to thank Stephen Ross, John Hartigan, Roger Ibbotson, Jonathan Ingersoll, Martin Shubik, Robert Shiller, Michael Montias, Ann Marie Logan, Burton Fredrickson, William S. Reese, Michael Tucker and the participants in the 1991 CRSP conference for helpful conversations and comments. I assume responsibility for all errors.

> Paintings are like stocks and a dealer is like a broker. Someone makes money, then there is someone else who's really good at investing in stocks, and he tells the investor what to buy. If someone tells you to go to a good gallery rather than one that's not so good, you'll get a painting that might turn out to be worth something, a painting you like that's also a good investment. Its like having a broker tell you what stocks to buy.
>
> Andy Warhol[1]

Introduction

The parallels that Pop master Andy Warhol drew between the stock market and the painting market provide an insider's insight into the issue of art market efficiency. If the art market were a perfect "informationally efficient" market, then Andy Warhol would have been wrong -- no art expert or gallery could help predict which paintings would appreciate in value and which would decline. In other words, if information about the future prospects of painting prices were common knowledge, then these prospects would already be reflected in *current* price levels. No dealer or gallery, no matter how clever, would know when the art market was poised for a boom or a bust, nor would they know which artist was "undervalued." The efficient market hypothesis has long dominated research in the stock market -- Warhol's observations to the contrary. While the efficient market hypothesis may logically apply to the stock market where differences of opinion about asset values play out in the continuous double-auction of the New York Stock Exchange every day, it is not clear that the hypothesis should apply to the art market, where purchase decisions are based on taste as much as upon discounted expectations of future price appreciation.

A logical consequence of the efficient market hypothesis as articulated by Fama (1970) is that asset prices should follow a random walk. In other words, past returns should not predict future returns. In fact, this implication is valid only to the extent that profitable timing strategies are allowed by transactions costs, and if superior returns may be obtained on alternative investments. In this paper, I examine the predictability of an index of art returns. I find that "weak form" efficiency, that is, predictions using past price levels, cannot be rejected for the art market. Despite this, there is some evidence of persistence in market returns, and consequently the profitability of timing strategies by knowledgeable market participants cannot be ruled out.

Another implication of an efficient market is that the immediate resale value of an asset will yield approximately what an investor paid for it. For instance, the purchaser of a share of stock on the NYSE should be able to immediately re-sell it for approximately the cost of the bid-ask spread, barring the release of new information, or a broad market move. This is a function of the high liquidity of the stock market, and the broad availability of financial information about share values. In this paper, I use transaction data to estimate the uncertainty surrounding an immediate resale of a painting. In contrast to the stock market, I find the re-sale uncertainty for art to be extraordinarily high. This large "price-risk" suggests that there may be great value to art dealers who, like specialists on the NYSE, can match buyers and sellers. When taken as a metric for art market efficiency, the price risk for paintings suggests that the art market has tended towards greater efficiency over its history. I attribute this trend to the increasing availability of auction information to prospective art buyers and sellers.

The paper is organized as follows. The next section describes the source of data. Section II describes the methodology for estimating the painting index used in the autocorrelation study and the methodology for estimating the price risk. Section III reports the results of autocorrelation tests. Section IV reports the estimates for the price risk of paintings. Section V concludes.

I. Data Sources

The empirical results about return predictability and price risk are based upon an analysis of a database of painting sales that extends over the period 1716 through 1986. From the three volumes of Gerald Reitlinger's *The Economics of Taste*, (1961,1963, 1971) and the annual volumes of Enrique Mayer's *International Auction Records*, (1971-1987) I identified purchase prices, sales prices, purchase dates and sales dates of 2,809 paintings that sold more than once. From this, I was able to form 3,329 compound returns, based upon the ratio of a painting's sale price to its purchase price.[2] Subsets of this data has been used by myself (Goetzmann, 1993b) and by previous authors (Anderson, 1974, Baumol, 1986, Frey and Pommerehne, 1989, 1990) to analyze the returns to investment in art, and to construct an index of art returns for comparison to the financial markets. While the four studies agree that art is a risky investment, Anderson (1974) and Goetzmann (1993b) find that the capital appreciation of art in the second half of the 20th century has been extraordinarily high.

II. Estimation Methodology and Empirical Model

Anderson (1974) and Goetzmann (1993) both construct indices of the capital appreciation of paintings using the "Repeat-Sales Regression." [RSR]. In this paper, I modified the procedure slightly, following Case and Shiller (1987), in order to estimate the price-risk component of art return variance.

To apply the RSR, we assume that the log of one plus the return for a certain asset *i* in period *t* may be represented by the log of one plus the return of an index of properties, μ_t, and an error term:

$$r_{i,t} = \mu_t + \epsilon_{i,t} \tag{1}$$

μ_t may be thought of as the average return in period t of properties in the portfolio. In order to estimate the vector μ. We wish to use sales data about individual properties to estimate the index μ over some interval t = 1....T. Our observed data consist of purchase and sales price pairs, $P_{i,b}$ and $P_{i,s}$ of the individual properties comprising the index, as well as the dates of purchase and sale, which we will designate with b_i and s_i . Thus,

$$r_i = \ln(\frac{P_{i,s}}{P_{i,b}}) = \sum_{t=b_i+1}^{s_i} r_{i,t} \qquad (2)$$
$$= \sum_{t=b_i+1}^{s_i} \mu_t + \sum_{t=b_i+1}^{s_i} \varepsilon_{i,t}$$

Case and Shiller (1987) show that this may be estimated by a three stage least squares regression procedure of the form:

$$\hat{\mu} = (X'\Omega^{-1}X)^{-1}X'\Omega^{-1}y \qquad (3)$$

Where X has a row of dummy variables for each asset in the sample, and a column for each holding interval, and the dummy takes on the value 1 if the asset is "held" in period t. In other words, when asset i is purchased in period b_i and sold in period s_i, X_{ij} takes in the value 1 for $b_i<j\leq s_i$, and zero otherwise, and Ω is a diagonal weighting matrix whose individual elements are the time between sales. In the first step, Case and Shiller set Ω equal to the identity matrix. They estimate the elements of Ω in the second stage of the procedure by regressing the squared errors from the OLS regression on an intercept and the time between sales:

$$\hat{e}_i^2 = \alpha+\beta(s_i-b_i)+\epsilon_i \qquad (4)$$

The slope from the regression estimate of (4) captures the effect of error variance compounding linearly through time, while the intercept term is a fixed quantity of risk unrelated to how long a work of art has been held. This fixed quantity is the component that Case and Shiller (1987) call the "price-risk." It captures the variance in price that is due to the purchase and sale transactions, rather than the dynamics of the price over the holding period. In the third stage, a GLS regression using the second stage estimates for the diagonal elements of Ω. Zeros are assigned to the off-diagonal elements.

In the following analysis, I use the RSR to estimate a decade-by-decade return index over the period 1720 through the partial decade ending in 1990. I use the price risk component estimated in (4) to quantify the informational efficiency of the art market.

III. Efficiency in Time Series: The Predictability of Art Returns

An efficient financial market is one in which asset prices rapidly and thoroughly impound all currently available information pertinent to the valuation of the assets in the market. There is an extensive literature examining the relationship between efficient market theory and the random walk model of stock market prices. In fact, a common first test of market efficiency is to examine the extent to which past stock returns predict future asset returns. This test of market efficiency by testing whether asset prices follow a random walk poses special problems in the art market, where paintings are traded infrequently.

It is difficult to test the random walk model with individual painting returns because a single painting will not sell frequently enough to test whether past returns predict future returns with any power. However, if we wish to use the portfolio of market returns estimated by the RSR we confront different problems. Goetzmann (1992) shows that the RSR regression typically induces spurious negative serial correlation in the estimated return series. Thus, a random walk test that measures the autocorrelation of returns would tend to reject market efficiency, and suggest that a contrarian strategy of buying when the market drops and selling when the market rises would be profitable. Positive serial correlation does not necessarily imply market inefficiency, however. A formal definition (see Fama, 1970) imposes the restriction that successive returns are independent and identically distributed. That is, $f(r_{t+1}|\Phi_t) = f(r_{t+1})$, where r_{t+1} is the return to a security next period, and Φ_t is the information set available today. In markets characterized by frequent trades, r_t is an element of Φ_t, and this further implies that successive returns are uncorrelated. In the case of the series $\mu_1 \ldots \mu_T$ estimated by the RSR, $\mu_1 \ldots \mu_t$ for $t<T$ is not in the information set Φ_{t+1}. As a matter of fact, in general, μ is only in the information set Φ_T. In simple terms, the RSR estimator $\hat{\mu}$ is forward-looking -- it incorporates information from the future into the current return estimate. Because of this, the RSR estimator is peculiarly ill-suited to random walk models of efficiency. Even if we found the market returns to be serially correlated, we could not conclude that the market is inefficient.

Of course, serial correlation may be of intrinsic interest. Case and Shiller (1987) note that the returns in the U.S. housing markets are persistent. They identify "bubbles" in market prices by examining the autocorrelation of successive appreciation returns. These are not necessarily due to irrational investor behavior, since we do not know the extent of the lag in the information set about past returns. In a sense, the natural temporal regularities of the housing market are not arbitraged away because the information may not exist to do so. We observe similar patterns in the art market.

I perform two tests of the random walk of painting returns. Each is designed to address particular problems posed by the RSR. The first test uses only "backward-looking" return estimates to predict the future. The RSR is performed for each decade using only sales information available to that point in time. The result of these "nested information set estimators" is a time series that we then use to predict the following decade return:

$$\mu^{t}_{NIS} = (X_t^{-1\prime}\Omega_t X_t) X_t^{\prime}\Omega_t^{-1}\rho_t \quad \textbf{(5)}$$

which we refer to as the nested information set estimator (NIS). X_t, , ρ_t, Ω_t are formed by using a subset of the data, defined by selecting only those observations for which the sales date is less than the date of the return to be estimated. The backward-looking restriction limits the data used in the RSR severely for the early periods. In fact, I am unable to estimate $r_{m,t}$ for decades before 1850. Perhaps because of this limitation, I find only marginal evidence of persistence in decade returns in Table I.

Table I

Regression of Bayes Decade Return Series on Preceding Decade "Nested Information Set" (NIS) Decade Return Series.

	Coefficient	t-Value
Intercept	.334	1.35
NIS Series	.493	1.70

R^2 = .19, N = 14, F = 2.905 on 1,12 df.

The second test of the random walk model of art market returns presents evidence for return persistence from one decade to the next, despite the fact that the market participants may be unaware of the pattern. As mentioned above, the RSR in general induces spurious negative autocorrelation in the estimated time series.[3] Case and Shiller (1987) address the negative autocorrelation problem by dividing the sample into two parts, estimating a return series with each half, and then regressing one of the series upon the lagged values of the other. If negative serial correlation is present, and the two series are highly correlated, this does not remedy the fundamental problem -- spurious negative or positive correlation in the lagged values across series is also a possibility.

To overcome this problem, we use a Monte-Carlo method to extend the Case and Shiller approach of dividing the sample. Instead of performing the test a single time, we perform it 100 times. In other words, for each simulation, we randomly group the sample into two halves, perform the RSR on each half and then regress one series on the lagged value of the other. Although we expect any single simulation to yield either spurious positive or spurious negative correlation, the distribution of the 100 regression coefficients should, on average be significantly greater than zero if return persistence is truly present in the sample. The results are reported in Table II.

Table II

Random Walk Test 2: Summary Statistics of Output For Two-Sample Monte-Carlo Regression of Lagged Series' (100 Samples)

Regression Statistic	Mean	Median	Standard Deviation
R^2	.16	.15	.10
Coefficient	.43	.43	.14
t-Statisic	1.96	1.92	.78

The average T-statistic of 1.96 is significant at the 95% level, and suggests strong evidence of persistence in art returns from decade to decade. Another way to test whether the coefficient of .43 is significantly different from zero is to divide by the standard deviation of the simulated coefficient distribution: .14. In effect, we have bootstrapped the regression coefficient, and the standard deviation of the distribution provides and estimate of the standard error of the coefficient. The bootstrap significance test indicates that the coefficient lies 3 standard deviations above zero. Considering the fact that the regression test is probably biased towards Type 2 error, these results are strongly indicative of the tendency of the market to move in trends.

One explanation for the persistance pattern in art returns is that they may be correlated to inflation, which is itself positively autocorrelated. This possibility is particularly pertinent to the issue of whether the the persistance in returns may be exploited. For instance, if the autoregression model predicts that art returns will fall, then an investor may not be able to exploit the decrease, because disinflation will cause the prices of alternative investments to fall as well.

The results of the NIS and the bootstrapping tests of autocorrelation in art returns suggests that the market displays some long-term serial dependency. This may be due to trends in inflation, or it may be due to slowly changing expected returns to art investment. Niether of these two possibilities may be exploited by trading rules, but, given the high transactions costs associated with art investment, it is unlikely that an investor would attempt a market timing strategy with art. None-the-less, the NIS autocorrelation test suggest that a dealer or market maker with access to past transactions information may be able to forecast the long-term future direction of the art market. Whether or not, as Andy Warhol suggested, this forecast can be exploited by judicious purchases and sales is an open question.

IV. Price Risk and Dealer Profits

Recall from the introduction that another implication of efficient market theory is that price information in the market is broadly available. Thus, the purchase price and the immediate resale price should not vary dramatically from each other. The price risk associated with immediate resale has been studied in the U.S. housing markets by Case and Shiller (1987) and by Goetzmann and Spiegel (1992). Both studies find that price risk for housing ranges between

5% and 10% of value. In other words, the price obtained in an immediate turn-around sale of a property may vary around the purchase price by 5 to 10%. Using the intercept term from the second stage of the repeat-sales procedure, we are able to estimate price risk for the art market. Under the assumption that half the price risk is due to the purchase and half is due to the sale, the intercept from (4) is multiplied by 1/2 and the square root is taken to estimate the price risk to which a seller is exposed. The estimate of price risk for painting sales is repoted in Table III. The estimate over the entire 271 year period suggests that the standard deviation of an immediate resale of a painting would be as much as 100%.

Table III

Changing Efficiency of the Art Market as Measured by Price Risk	
Sub-Set	Single-Sale Standard Deviation
Entire Sample	104.0 %
Bought Before 1800	70.7 %
Bought Before 1850	66.9 %
Bought Before 1900	61.2 %
Bought Before 1950	51.4 %

Table III also reports the price risk for pictures purchased at later and later dates. Notice that the price risk spread has decreased steadily through time. It is unlikely that this pattern is due to errors in the exchange rate estimates, since exchange rates were fixed for most of the pre-20th century prices. By conditioning on our measure of the price risk upon the purchase date of the painting, we find evidence that, even though the basic institutions of the art market may have evolved as early as the 18th century, it appears to have steadily tended towards greater efficiency as a mechanism for processing information and matching buyers and sellers. Part of this increase in efficiency may be due to the increase in quality and quantity of public information about auction prices. Comprehensive international at auction records have been available in printed form since 1970. Since the mid-1980's, interested buyers and sellers have been able to access online databases such as ArtQuest, Inc. to search for auction transactions by artist, medium, size and subject matter of painting. This type of information will inevitably reduce the uncertainty about the current value of a work of art. Another source of increase in efficiency is the development of art historical research. Despite a long history of active art scholarship, It has only been in the twentieth century that many painters have *catalogue raisoneés*, or comprehensive accountings of their *ouevre*. Art histocial publications "level the field" for woners and prospective buyers of art, by identifying the range of styles, characteristic signatures and major collections of an artist. Neither good auction information, nor a good library are substitutes for market experience. A dealer specializing in the work of a single artist will information about the quantity and location of the artist's works, as well as about the prices for which that work has sold. In addition, a dealer will also know who is in the market for a particular artist's work. Until the price risk of owning a painting drops to much narrower range, this market information will continue to provide substantial rewards.

The large spread between the "true" value and the market value of a painting has broader implications beyond the evolution towards efficiency of the art market. Even if one were unaware of the extensive network of dealers and valuation experts in the art market, the results of our price risk analysis would suggest their existence. Dealers profit by having superior information about the paintings, or about the collectors interested in the paintings. A gap of 50% in the going rate for a particular work of art suggests that profits may high for a market-maker. The high price risk represents a considerable motivation for a dealer to develop superior information about the assets, and about the market demand for them. For example, if the price variation in the auction market were similar in magnitude to the price variation in the dealer market, then an art dealer who is attempting to purchase low and sell high might expect to mark prices up by 50%.

V. Conclusion

The art market is considerably less liquid that the financial markets. This illiquidity may be a cause or an effect of the mechanisms by which information is processed by the structure of collectors, dealers, agents and auction houses that currently make up the art market. In this paper, I have sought to determine the extent to which broad trends in the art market may be predictable by past transactions information. The answer appears to be weakly affirmative. I have also sought to quantify informational inefficiency through a measure of "price risk." Estimates indicate that the art market has five to ten times more price risk than the U.S. housing market. This is bad news for art investors seeking to liquidate their properties at a profit. They appear to expose themselves to more than 50% standard deviation in the realized price when they take a painting to auction. This transaction risk implies that there are large potential rewards to art dealers who can match buyers and sellers. The 50% price risk estimate may translate into a potential 50% dealer markup. The emergence of information vendors and the growth of art historical scholarship may have contributed to the decrease in price risk over the history of the art market. If so, then I would anticipate such a trend to continue.

References

Anderson, Robert C., 1974, "Paintings as Investment," *Economic Inquiry*, March, 2, 13-26.

Baumol, William, 1986, "Unnatural Value: or Arrt Investment as a Floating Crap Game," *American Economic Review*, May (Papers and Proceedings) 76, 10-14.

Bidwell, R.L., 1970, *Currency Conversion Tables; A Hundred Years of Change* London, Rex Collins.

Case, Karl and Shiller, Robert, 1987, "Prices of Single Family Homes Since 1970," *New England Economic Review*, September-October, 45-56.

Collins, Michael, 1986, "Sterling Exchange Rates, 1847-80," *Journal of European Economic History*, September, 15, 511-533.

Fama, Eugene, 1970, "Efficient Market: A Review of Theory and Empirical Work," *Journal of Finance* 25, 2, 1181-1185.

Frey, Brono and Pommerehne, Werner, 1989, *Muses and Markets: Explorations in the Economics of the Arts*, London, Blackwell.

Frey, Brono and Pommerehne, Werner, 1990, "Art Investment: An Empirical Inquiry," *Southern Economic Journal*, 56, 396-407.

Goetzmann, William N., 1992, "The Accuracy of Real Estate Indices: Repeat Sales Estimators," *Journal of Real Estate Finance and Economics*, 5, 5-53.

Goetzmann, William N., 1993a, "Patterns in Three Centuries of Stock Market Prices," *Journal of Business*, 66,2, 249-270.

Goetzmann, William N., 1993b, "Accounting For Taste: Art and the Financial Markets Over Three Centuries," *American Economic Review*, December, 83, 4.

Goetzmann, William N. and Spiegel, Matthew, "Temporal and Non-Temporal Components of Residential Real Estate Price Appreciation," Working Paper, Columbia University.

Lo, Andrew, 1991, "Long-Term Memory in Stock Prices," *Econometrica*, 59, 1279-1314.

Madelbrot, Benoit, 1972, ""Statistical Methodology For Non-Periodic Cycles: Form the Covariance to R/S Analysis," *Annals of Economic and Social Measurement*, 1, 259-90.

Mayer, Enrique, 1971-1987, *International Auction Records*, New York, Archer Fields.

Ratcliff, Carter, 1983, *Andy Warhol*, New York, Abbeville Press.

Reitlinger, Gerald, 1961,1963,1971, *The Economics of Taste*, London, Barrie and Rockcliff.

Notes

1. Quoted in Ratcliff (1983) from an interview with the artists in 1980.

2. The database may suffer from a variety of biases. Paintings that fell out of fashion, as well as paintings that ended up in the permanent collection of museums would not have been sold twice. Some of the transactions prices are bought in prices. Most but not all of the records are auction records, and participants in the art market regard auctions are "riskier" than dealer sales. All prices were converted to British currency using R.L. Bidwell (1970), Michael Collins (1986) and *The International Statstics Yearbook.* End-of-year conversion ratios inevitably add noise to the data.

3. Alternative test statistics such as the rescaled range suggested by Mandelbrot (1972), modified by Lo (1991) and applied to the long-term stock price series' by Goetzmann (1993b), are likewise biased. The rescaled range will be biased downwards due to spurious negative autocorrelation suggestive of mean-reversion. Goetzmann (1993b) shows how the rescaled range bias may be adjusted via bootstrappping procedures. The rescaled range for the decade index of art returns is 1.54, suggesting returns are persistant, rather than reverting.