

The art market: what do we know about returns?

Ventura Charlin (1) and Arturo Cifuentes (2)

- (1) V.C. Consultants, Los Leones 1300, Suite 1202, Santiago, CHILE,
e-mail: ventcusa@gmail.com
- (2) Financial Regulation Center, Faculty of Economics and Business, University of Chile, Santiago, CHILE,
e-mail: arturo.cifuentes@fen.uchile.cl

KEY WORDS: hedonic price indices; hedonic regression; wild bootstrap; bootstrap method art market; art market returns; hedonic models; alternative investments; bootstrapping

JEL CODES: Z11, Z10, G11, C01, C02

Abstract

We examine the annual returns based on auction data for two groups of artists (Surrealists and Impressionists) and two individual artists (Picasso and Renoir) using hedonic pricing models in combination with a wild bootstrap statistical technique. This approach allows us to estimate confidence intervals for such returns. In addition, we estimate confidence intervals for several figures of merit based on these returns, such as correlations with returns of other type of assets, and risk-return metrics.

We find that the confidence intervals associated with these figures of merit are so wide that it is difficult, if not impossible, to derive absolute conclusions or to make meaningful comparisons, with the behavior of other assets. We also observe that relying on single-point estimates of the above-mentioned metrics –without accounting for the corresponding confidence intervals– can lead to erroneous interpretations regarding art-market returns.

Moreover, our results suggest that previous studies regarding art market returns, their correlation with broader market indices, and their risk-return profiles, should be re-examined as they were based on single-point estimates of the relevant metrics. Finally, these findings might be of interest to researchers who use hedonic pricing models in the analysis of other infrequently traded assets as the number of sales/observations is likely to be rather low.

1 Introduction

Paintings, notwithstanding their artistic merits, are valuable financial assets. Consequently, many authors have studied the returns of the art market, as well as their correlation with the returns of more conventional assets such as stocks and bonds, in order to assess the potential benefits of investing in art vis-à-vis such alternatives.

A key challenge to study returns in this market is that the relevant prices are not as clearly defined as in other markets. Consider stocks for example: the return r on a stock between t_1 and t_2 , is simply $(P_2/P_1) - 1$, where P_1 and P_2 , the corresponding stock prices, are well-defined quantities. A return computed with such expression is the real (true or actual) return; it is not an estimate, for P_1 and P_2 are known quantities: observed prices at the times of interest. Hence, there is no error associated with r .

However, if we wish to compute the return on Picasso paintings between t_1 and t_2 , the calculation is not that straightforward for P_1 and P_2 are not clearly defined (unless we are referring to a specific painting that has been sold twice, and at precisely those two times). Picasso paintings, despite the fact that come from the same artist, are essentially different objects: they may be similar, but they are not identical. Therefore, to compute r we need to rely on representative prices at t_1 and t_2 that should somehow capture the existing price information regarding what are actually a group of heterogeneous objects. Evidently, these representative prices are just proxies for P_1 and P_2 , and as such, will have an error associated with them.

Hence, several techniques have been developed to estimate representative prices, and in turn, returns. It should be noted, however, that the r obtained in these cases, unlike the return computed in the case of stocks, is only an approximation to some elusive ideal return. Therefore, r should be characterized in terms of a point estimate or average, plus a confidence interval. The first issue (how to estimate P_1 and P_2) has been treated in detail in the academic literature, as we will see in the next section. The second point, however, has been ostensibly neglected in previous art market research. That is, previous studies about the art market have based their conclusions on calculations carried out with single-point estimates of r , disregarding the uncertainty around those estimates.

In short: in the art market there is no such thing as a correct or true return; thus, we need to acknowledge that we are dealing with an estimate of what we hope is a representative return and consequently there is a margin of error involved. Therefore, any computation we carry out based

on those return estimates must incorporate some measure of their dispersion, and should not be based only on a single-point estimate. Otherwise, it is impossible to assess the relevance of the figure of merit computed.

2 Previous work

2.1 Returns

Anderson (1974) appears to be the first researcher who explored the art market from an econometric viewpoint. Using sales data for the period 1643-1970 he concluded that paintings had offered a return that was about fifty percent the return offered by common stocks. He also suggested that sales of paintings in recent years had done better in terms of approaching the returns offered by stocks, but emphasized that there were important differences depending on the artists and the type of school or movement.

Baumol (1986), in other widely cited paper, challenged the notion that there could be anything approaching an equilibrium price in the market for paintings. He contrasted the market for paintings with the stock market and highlighted that in the case of stocks the "true" or equilibrium price is known whereas in the art market such concept is imprecise. Finally, analyzing data from 640 painting transactions between 1652 and 1961, he concluded that returns offered by paintings did not compare favorably with those of government securities and exhibited a remarkable degree of variability. His main conclusion (captured succinctly by his paper's title "Art Investment as a Floating Crap Game") was that if predicting stock prices was difficult, predicting art market prices was probably a hopeless task.

More recently, a number of authors have expressed, with different degree of emphasis and with different caveats, a somewhat uniform verdict: risk-adjusted returns associated with paintings do not appear attractive when compared with stocks and bonds. For instance, Renneboog and Spaenjers (2013), using data covering the period 1957-2007, built an art index that exhibited a modest 3.97% real annual return expressed in U.S. dollars. That is, a performance similar to that of corporate bonds but with much higher risk. Mandel (2009) also found similar results for the period 1950-1999, namely, that art exhibited returns lower than both the S&P 500 and the Dow Jones industrial index, but with higher volatility. Worthington and Higgs (2004), using data of some specific art market segments (old masters, surrealists, impressionists, 19th century European, etc.) found that returns on paintings showed lower rates and higher volatility when compared with more conventional assets. Other authors that have arrived at similar conclusions

are Renneboog and Van Houte (2002), Ashenfelter and Graddy (2003) and Agnello (2002). Campbell (2008) summarizes art market returns obtained by different authors for different segments and time-periods. A comprehensive literature review regarding returns can be found in the paper by Renneboog and Spaenjers (2013).

2.2 Correlation

The findings regarding correlation between art returns and the returns experienced by more traditional assets are more ambiguous. For example, Mandel (2009), as well as Goetzmann (1993) and Stein (1977) before, found that the art market and equities were highly correlated. This is in contrast with Mei and Moses (2002) who concluded that paintings were lowly correlated with equities (S&P 500), as well as treasuries and corporate bonds. Campbell (2008) also found low correlation between art returns in general and both, stocks (MSCI world stock index), and bonds (treasuries and corporates). Worthington and Higgs (2004) detected a 16% correlation between art and large-company stocks and a -31% correlation between art and small caps. Renneboog and Spaenjers (2013) reported a -3% correlation between art and the S&P 500 index, a 20% value with respect to global stocks, and somewhat higher values (in the 30%-45% range) when referencing commodities and real estate indexes.

Clearly, some of these discrepancies can be attributed, in principle, to having examined different time-periods, different art market segments, and the use of different techniques to estimate returns. However, it is safe to say that it is difficult to make sweeping statements regarding the correlation between the art market (or some segments of it) and broader market indexes.

2.3 Estimating returns in the art market

The difficulties associated with estimating returns in the art market are the result of dealing with groups of heterogeneous goods. There are two established methods to estimate price variations (and therefore returns) in such situations: (i) repeat-sales regressions (RSRs); and (ii) hedonic pricing models (HPMs).

RSR-methods estimate returns based on price information regarding paintings that have been sold at least twice (e.g., Anderson 1974; Baumol 1986; Goetzmann 1993). This way the problem of dealing with heterogeneous goods is circumvented. Nevertheless, this method has two disadvantages: a potential selection bias and the fact that it only employs a small subset of the available information, typically, below 25%. This is a characteristic of the art market. In the

U.S. real estate market, for example, where the Case-Shiller index (a RSR-based index) is the standard, repeat sales account for more than 90% of total sales.

HPMs –the approach favored by most researchers and the method on which the bulk of previous conclusions regarding the art market are based– are better suited to manage product variety and have the advantage that they use all the available data (e.g., Chanel et al. 1994, 1996; de la Barre et al. 1994; Edwards 2004; Renneboog and Spaenjers 2011, 2013; Renneboog and Van Houtte 2002). The idea behind this method is to employ a regression in which the natural logarithm of a painting’s selling price is the dependent variable. The independent variables are linear or higher-order polynomial expressions based on the age of the artist at the time the painting was executed and variables associated with the characteristics of the painting and auction-related information, plus a set of time-dummies linked to the year the painting was sold. The general HPM characterization is discussed in more detail later in the paper.

Additionally, a few studies have also explored the benefits of using simpler metrics such as the geometric mean of the prices observed on a given year (Chanel et al. 1996; Worthington and Higgs 2004). More recently, Charlin and Cifuentes (2014) proposed another approach to estimate returns based on the price per unit of area or APV (a normalized price that homogenizes the paintings sales data by controlling for the area). This metric, which has several appealing features beyond its simplicity, has not yet been widely used in return-related studies.

In summary, previous studies (most of them based on HPMs) have concluded that the art market does not offer attractive risk-adjusted returns when compared with more conventional assets, while conclusions regarding the correlation between art and other markets are somewhat less clear. We should note that these findings were based on point estimates of returns; whether these conclusions would hold if the confidence intervals of such estimates had been considered is unknown.

A handful of authors have computed average returns by simply averaging the year-to-year return estimates based on the HPM-regression coefficients and then they calculate the corresponding standard deviation of those returns in an attempt to have a sense of the dispersion associated with such average. This approach is not satisfactory for it fails to address the fact that those estimates are biased as a result of the log-transformation (a subject treated in a subsequent section). As far as we know, Renneboog and Spaenjers (2013) are the only authors who have corrected their point return estimates for such bias. In that sense, their computations marked an improvement compared to previous studies. However, they failed to calculate confidence intervals associated

with their estimates. We should also mention that Chanel et al. (1996), in a study to investigate the merits of RSRs versus conventional HPMs, estimated confidence intervals for the time dummy coefficients of the HPM regression. Nevertheless, they did not perform the additional step of employing these intervals to assess how they impacted the accuracy of the return estimates.

With that background the goals of our paper are: (i) assess how the time-dummy coefficients' standard errors influence the errors in the art return estimates when using hedonic models; (ii) then, in turn, assess how the errors in such returns translate into errors in the correlation estimates between the art market and other assets (namely, the U.S. stock market as described by the S&P 500 index); and (iii) demonstrate the effectiveness of a wild bootstrap technique to estimate confidence intervals for several figures of merit.

3 Data

We employ four data sets in this study:

- a. Data set A consists of 3,780 observations of paintings auction prices covering the period [March 1988 – December 2014]. This data set gathers information from five artists: Giorgio de Chirico, Joan Miro, Max Ernst, Rene Magritte, and Roberto Matta. We refer to this group as the surrealists.
- b. Data set B consists of 2,121 observations covering the period [March 1985 – December 2012]. This data set gathers information from six artists: Alfred Sisley, Camille Pissarro, Claude Monet, Odilon Redon, Paul Gauguin, and Paul Signac. We refer to this group as the impressionists.
- c. Data set C consists of 1,972 observations of Pierre-Auguste Renoir's paintings auction prices and their characteristics covering the period [March 1985 – December 2014].
- d. Data set D consists of 1,322 observations of Pablo Picasso's paintings auction prices and their characteristics covering the period [March 1985 – December 2014].

The databases were built based on information provided by artnet (www.artnet.com) and supplemented by auction data provided by the Blouin Artinfo website (www.artinfo.com).

We adjusted all prices to January 2010 U.S. dollars (using the U.S. CPI index) and expressed them in terms of premium prices (we modified hammer prices and expressed them in terms of

equivalent premium prices when appropriate). In addition, we eliminated all observations where the selling price was below US\$ 10,000 or less than 1 US\$/cm². Sotheby's and Christie's dominate the data sets, as together they account for 77% of the sales.

The selection of painters was somewhat arbitrary and with no consideration to artistic merits. The idea was to have data sets pertaining to individual artists as well as group of artists. The chief consideration was to select artists that had enough information over a reasonable time-period.

Table 1 summarizes the key features of the four data sets. Table 2 includes some additional information specific to the artists included in data sets A and B.

4 Method of analysis

4.1 General framework

Let P_i^t be the observed auction price of the i_{th} painting at time t . We consider a general HPM, whose structure can be described as follows:

$$\text{Log}(P_i^t) = \alpha_0 + \sum_{n=1}^S (\alpha_n Y_i^n) + \sum_{n=S+1}^N (\alpha_n Z_i^{n,t}) + \sum_{m=1}^M (\beta_m T_i^m) + \varepsilon_i^t \quad (1)$$

where Y_i^n ($n=1, \dots, S$) denotes the value of the n^{th} time-independent characteristic of painting i (for example, the size of the painting or the age of the artist when the painting was created); $Z_i^{n,t}$ ($n = S+1, \dots, N$) denotes the value of the n^{th} time-dependent characteristic of painting i (for example, the auction house where the sale took place at time t or whether it was sold in an evening or a day sale); T_i^m is a time-dummy that takes the value 1 when $m = t$ (otherwise is 0); and ε_i^t is the error term. The α 's and β 's are the regression coefficients. We have N characteristics and M time periods. The return between two consecutive periods, v and $v+1$, is generally estimated as $(P_{v+1}/P_v) - 1$ where $P_v = \exp(\beta_v)$.

Notwithstanding their popularity, HPMs exhibit a shortcoming: even though the estimates of the β 's provided by the HPMs are unbiased, the return estimates are not. This is a direct consequence of the log-transformation. This topic has been discussed in detail by Dalen and Bode (2004) in the context of HPMs applied to price indexes. The existence of this bias, in more general settings, had been detected and analyzed before (Goldberger 1968; Teekens and Koerts 1972).

Accordingly, our analyses proceeded in three steps:

(1) We estimated, using a HPM, point estimates of the year-to-year returns based on the β coefficients resulting from the HPM regression, as well as some derived quantities from such estimates, namely, the risk profile of such returns and their correlation with the S&P 500. We define the risk profile (RR) of returns as the standard deviation of such returns divided by their mean value.

(2) We recalculated the figures of merit described in (1) based now on point return estimates corrected for the log-transformation bias using the correction employed by Renneboog and Spaenjers (2013) and detailed by Triplett (2004) and Silver and Heravi (2007).

The correction works as follows: the conventional (biased) year-to-year return estimate between two consecutive periods, based on the HPM time-dummy coefficients, is $(P_{v+1}/P_v) - 1$ where $P_v = \exp(\beta_v)$. The bias correction consists of replacing $\exp(\beta_v)$ by $\exp(\beta_v^C)$ where $\exp(\beta_v^C) = \exp[\beta_v + (1/2)(\sigma_v^2 - \sigma_0^2)]$ and σ_0^2 and σ_v^2 represent the variances of the residuals of price observations at periods 0 and v, respectively.

(3) We used a wild bootstrap statistical technique in combination with the corrected point return estimates to determine both, single-point estimates, plus their corresponding confidence intervals, for all the relevant figures of merit.

4.2 Wild bootstrapping resampling method

Bootstrapping refers to taking many samples with replacement from the original data to create a large number of plausible data sets (Efron and Tibshirani 1993). This way, we can improve the accuracy of the initial return point estimates derived from the regression time-dummies (in this case, corrected for the bias) by averaging from these simulated samples. The confidence intervals are obtained non-parametrically by determining, in this case, the 2.5% and 97.5% cutoff points from the resulting distributions.

There are several ways to carry out bootstrapping simulations in the context of regressions. The three most common methods are known as case-based resampling, model-based resampling, and wild bootstrap (Liu 1988).

In this study we used a wild bootstrap technique since it has the advantage of being more suitable to overcome the difficulty presented by heteroskedastic errors. This method was developed by Liu (1988) and further advanced by Davidson and Flachaire (2008). The wild bootstrap is similar

to the model-based resampling method except for one difference: in the model-based approach, the predictors in the regression are treated as fixed, and therefore resampling is done from the residuals generated by the regression based on the original sample. And then, for each observation, a new price is created by adding the selected residual to the price predicted by the regression model. These issues have been discussed in detail by Beer (2007), Dalen and Bode (2004), and Stine (1989). In Beer (2007) it is recommended that the raw residuals $\hat{\xi}_i$ be modified following Davison and Hinkley's (1997) formula to make sure their variances matched those of the random errors. Thus, the modified residuals ξ_i^* are:

$$\xi_i^* = \frac{\hat{\xi}_i}{(1 - \widehat{\text{leverage}}_i)^{1/2}}$$

where $\widehat{\text{leverage}}_i$ is the i^{th} diagonal element of the hat matrix $\widehat{\mathbf{H}}^t$ of the regression model at time t . Thus, $\widehat{\mathbf{H}}^t = \mathbf{X}^t(\mathbf{X}^{t'}\mathbf{X}^t)^{-1}\mathbf{X}^t$, where X is the design matrix.

In the wild bootstrap method a difference with the conventional model-based approach is introduced: the residuals ξ_i^* , before being added to the values predicted by the regression to create the new sample, are modified again, this time, using a random number drawn from the Rademacher distribution.

The Rademacher distribution is defined as

$$f(x) = \begin{cases} -1, & \text{with probability 0.5} \\ 1, & \text{with probability 0.5} \end{cases}$$

Thus, we multiply the modified residual by a discrete random variable having the value 1 or -1 with equal probability. And such re-modified residual is then added to the value predicted by the regression to generate a new price-observation. This method is explained in detail in Beer (2007). Its main advantage is that it has the ability to incorporate heteroscedastic error terms, and according to Flachaire (2005) the wild bootstrap performs better than the case-based bootstrap.

We generated 2,000 wild bootstrap replications. This number of replications was deemed appropriate based on the marginal improvements observed with 5,000 replications.

5 Results and discussion

The results were fairly consistent for all four data sets, thus, we present in detail the results based on the surrealists, and in summary fashion those of the other three data sets.

5.1 The surrealists under the eye of the HPM

The HPM fitted had an adjusted $R^2=0.72$. The most relevant characteristics, with significant α 's using HC2 as the heteroskedasticity consistent covariance matrix estimator (HCCME) that was proposed in MacKinnon and White (1985), were: whether the painting was painted by Matta ($t=-35.11$, $p<.0001$); the natural logarithm of the area of the painting ($t=20.59$, $p<.0001$), whether the painting was painted by Magritte ($t=15.17$, $p<.0001$); and whether the painting was sold in an evening sale ($t=14.30$, $p<.0001$).

Table 3 displays the results. The first five columns correspond to the information typically presented in most art market analyses. The sixth column shows the estimated year-to-year return after applying the bias correction. It is obvious that the log-transformation bias has an important effect and cannot be ignored. As mentioned before, most previous analyses only report the single-point estimate without applying the bias correction, with Renneboog and Spaenjers (2013) being the only exception.

A far more important observation refers to the β 's standard errors: they are not small. This should serve as a warning regarding the accuracy of the single-point return estimates. Curiously, even though some researchers do report the standard errors, they do not use them for any further calculation or accuracy assessment. We suspect this is due to the fact that there is no analytical expression to link the β 's standard errors and some dispersion-related metric associated to the return estimates. This takes us to the bootstrap results reported on columns 7 through 10. Two interesting features: (i) the necessity of employing the bootstrap technique (or other suitable approach) to improve the accuracy of the return single-point estimates becomes apparent when we compare the results reported in columns 6 and 7. The discrepancy between these two columns is noticeable. In other words, the bootstrap estimate represents an important improvement in comparison with the estimate provided solely based on the HPM; and (ii) the widths of the return confidence intervals are so wide (columns 9 and 10) that is almost impossible to even conclude if the return for some periods was positive or negative. This situation shows clearly the danger of attempting to derive any conclusion from analyses based only on single-point estimates of returns.

Table 4 reports the summary metrics for the surrealist artists for the 1988-2014 period. Again, and in agreement with the trends shown in Table 3, the confidence intervals estimated by the bootstrap statistical technique, for all metrics, are sizable. It is also clear that any conclusion regarding returns based on the 2.48% value (the single-point estimate before applying the bias correction) even if one computes the corresponding standard deviation, gives a very incomplete if not misleading picture in terms of the returns. The fact that the risk-return (RR) metric shows less variability is probably attributable to the fact that the errors in the numerator and denominator have the same sign and their effects tend to cancel out.

The existence of a bias when one estimates returns through the HPM's β coefficients has been reported by previous authors in other applications not related to the art market; see, for instance, Goldberger (1968), Triplett (2002), and Dalen and Bode (2004). It has been observed that it is more pronounced when dealing with price changes over a long period of time, say, one-year intervals, than shorter periods (e.g., hours, days, weeks, etc.) In many applications, such as month-to-month inflation indexes based on pooled data this consideration is rather academic, as most characteristics do not change dramatically between two consecutive periods and the bias is negligible. However, in art market studies, where most research is conducted in one-year periods, and this bias is usually associated with violent changes in both characteristics and prices, ignoring this bias might be ill-advised, as our results indicate.

Briefly, the results underscore that it is not possible to derive any conclusions from a simple hedonic regression without further refining the initial estimates via a bootstrapping method –or other suitable algorithm– since the effects of the regression coefficients' standard errors in the subsequent calculations are huge and cannot be ignored.

Unfortunately, the majority of the conclusions reported in the literature regarding art returns and correlations are derived from computations based on point estimates of the β coefficients; they rely on the type of information displayed in the third column of Table 3. See for instance, Agnello (2006), Bakhouché and Thebault (2011), Campos and Barbosa (2009), Kraeussl and Lee (2010), Lazzaro (2006), Lucinska (2015), Onofri (2009), Pesando and Shum (1999), Renneboog and Spaenjers (2013), and Sproule and Valsan (2006). Note that some authors such as Renneboog and Spaenjers (2013), or Pesando and Shum (1999), have reported the standard errors associated with the time-dummy regression coefficients--and in the case of Renneboog and Spaenjers (2013) they have even attempted to correct for the bias in the return estimates--however, they have failed to estimate confidence intervals for the average returns or the

correlations with other assets. Thus, the validity of any conclusions stemming from these findings should be taken with a fair amount of skepticism.

5.2 All four data sets

Table 5 summarizes the overall results regarding the four data sets. The results reinforce the tendencies already described in Tables 3 and 4, namely: (i) the log-transformation bias has an important effect on the single-point estimates; (ii) there are significant differences between the single-point (bias-corrected) estimates and the corresponding figures reported by the bootstrap algorithm; and (iii) the confidence intervals for the annual returns and correlations are huge. The RR metric, as before, shows less variability.

The results regarding the returns are unsettling. Take Renoir for example: in spite of having 1,972 observations (a huge number by all accounts, especially when dealing with a single artist) we can only say that the average return for the period considered was between 5.90% and 10.30%. What can we say regarding the performance of Renoir as an asset class, vis-à-vis the performance of an index whose year-to-year return can be computed with no error? Not much, unfortunately.

We can speculate that for databases with many more observations, perhaps--and this is only a hope--the confidence intervals might be narrower. However, most investors are interested in assessing the performance of an individual artist--or at most, a distinctive school of artists--and therefore, the corresponding data sets will be comparable in sizes to what we have in this study.

The results regarding the correlation between the artists considered and the S&P 500 are even more disturbing (and in part they might explain the lack of consistency in terms of the findings by previous studies). In two of the four cases (Impressionists and Picasso) we are not even sure of the correlation sign (the confidence interval includes the zero). Under these circumstances it is difficult to say if adding these assets to a portfolio of stocks could add or not to its diversity or can help to mitigate its overall risk. We are not aware of any previous study that had looked at the correlation between art and stocks considering the error in the correlation estimates.

6 Conclusions

Art return estimates, and other metrics based on the β coefficients of HPMs, should be taken with a great deal of care. The errors in the β 's are likely to be large, and therefore, have an important effect on the accuracy of these metrics. More precisely, not knowing what the confidence

intervals associated with such metrics are, renders them almost worthless. Unfortunately, all previous art-related studies have based their conclusions on single-point estimates of the relevant metrics disregarding their confidence intervals.

The good news is that the bootstrapping technique demonstrated here is a useful tool to estimate such confidence intervals. The bad news is that the resulting confidence intervals can be quite wide. In other words, our results show that it is almost impossible to make even mildly reliable –or interesting– statements regarding the returns offered by art investments or the correlation between such returns and the returns of more conventional assets such as stocks and bonds.

It is important to realize that studies dealing with infrequently traded assets such as paintings (especially if the analysis refers to only one artist as opposed to a group of artists) and other collectibles (stamps, musical instruments, or classic cars, for instance), are likely to have a limited number of observations. In fact, in all probability, fewer observations than in the examples discussed herein. In these instances, the importance of estimating the confidence intervals of the relevant figures of merit is even more critical as the errors associated with the single-point estimates are likely to be larger.

Additionally, a number of authors have applied the CAPM-model in the context of art returns (for instance, Stein 1977; Chanel et al., 1994; Hodgson and Vorkink 2003; Kraeusl and Lee 2010). The fact that they performed their computations without taking into account how the errors in the return estimates impacted the accuracy of the CAPM-model coefficients should be taken as a warning against accepting their conclusions at face value.

Finally, back to Baumol (1986) and his landmark pronouncement: "Art Investment as a Floating Crap Game." Ironically, almost thirty years later, it is still one of the few absolutes in the art market. When all is said and done, making definite statements about art returns might well prove to be as contentious as deciding who the best artist is. The more things change, the more they remain the same.

References

- Agnello, J. (2002) Investment Returns and Risk for Art: Evidence from Auctions of American Painters, *Eastern Economic Journal*, 28(4), 443–463.
- Agnello, R. (2006) Do U.S. Paintings Follow the CAPM? Findings Disaggregated by Subject, Artist, and Value of the Work, available from <https://ideas.repec.org/p/dlw/wpaper/06-02.html>

- Anderson, R. C. (1974) Paintings as an investment, *Economic Inquiry*, 12, 13–26.
- Ashenfelter, O. and Graddy, K. (2003) Auctions and the price of art, *Journal of Economic Literature*, 41, 763–86.
- Bakhouché, A. and Thebault, L. (2011) What Determines Cezanne's Art Pricing, A Hedonic Regression Method, Scientific Annals of the "Alexandru Ioan Cuza" University of Iasi, Economic Sciences Section, available from <http://anale.feaa.uaic.ro/anale/en/Arhiva202011%20BakhouchéThebault/421>
- Baumol, W. J. (1986) Unnatural value: Or art investment as a floating crap game, *American Economic Review*, 76, 10–14.
- Beer, M. (2007) Bootstrapping a Hedonic Price Index: Experience from Used Car Data, *Advances in Statistical Analysis*, 91, 77–92.
- Campbell, R. (2008) Art as a Financial Investment, *Journal of Alternative Investments*, 10(4), 64–81.
- Campos, N. and Barbosa, R.L. (2009) Paintings and Numbers: an Econometric Investigation of Sales, Prices, and Returns in Latin American Art Auctions, *Oxford Economic Papers*, 61, 28–51.
- Charlin, V. and Cifuentes, A. (2014) An Investor-Oriented Metric for the Art Market, *Journal of Alternative Investments*, 17(1), 87–101.
- Chanel, O., Gerard-Varet, L. A., and Ginsburgh, V. (1994) Prices and returns on paintings: An exercise on how to price the priceless, *The Geneva Risk and Insurance Review*, 19, 7–21.
- Chanel, O., Gerard-Varet, L.A., and Ginsburgh, V. (1996) The Relevance of Hedonic Price Indices: The Case of Paintings, *The Journal of Cultural Economics*, 20, 1–24.
- Dalen J. and Bode, B. (2004) Estimation Biases in Quality-Adjusted Hedonic Price Indices, Rotterdam School of Management, Erasmus University, Rotterdam, available from www.ipeer.ca/papers/vanDalenBodeOct.1,2004,SSHRC%20Paper17.pdf
- Davidson, R. and E. Flachaire (2008). The wild bootstrap, tamed at last, *Journal of Econometrics*, 146 (1), 162–169.
- Davison, R. and Hinkley, D.V. (1997) *Bootstrap Methods and Their Applications*, Cambridge University Press, Cambridge.
- de la Barre, M., Docclo, S., and Ginsburgh, V. (1994) Returns of impressionist, modern and contemporary European paintings 1962–1991, *Annales d'Economie et de Statistique*, 35, 143–81.
- Edwards, S. (2004) The Economics of Latin American Art: Creativity Patterns and Rates of Return, *Economía*, 4, 1–35.
- Efron, B. and Tibshirani, R.J. (1993) *An Introduction to the Bootstrap*, Chapman and Hall, New York.
- Flachaire, E. (2005) Bootstrapping heteroskedastic regression models: wild bootstrap vs. pairs bootstrap, *Computational Statistics & Data Analysis*, 49 (2), 361–376.

- Goetzmann, W.N. (1993) Accounting for taste: art and financial markets over three centuries, *American Economic Review*, 83, 1370–76.
- Goldberger, A.S. (1968) The Interpretation and Estimation of Cobb-Douglas Functions, *Econometrica*, 36(3/4), 464–472.
- Hodgson, D. and Vorkink, K. (2004) Asset pricing theory and the valuation Canadian paintings. *The Canadian Journal of Economics*, 37(3), 629-655.
- Kraeussl. R. and Lee, J. (2010) Art as an Investment: The Top 500 Artists, Department of Finance, VU University, Amsterdam, available from <http://www.efmaefm.org/0EFMSYMPOSIUM/2011-Toronto/papers/Kraeussl.pdf>
- Lazzaro E. (2006) Assessing quality in cultural goods: The hedonic value of originality in Rembrandt's prints, *Journal of Cultural Economics*, 30(1), 15-40.
- Liu, R. Y. (1988). Bootstrap procedures under some non-I.I.D. models, *Annals of Statistics*, 16, 1696–1708.
- Lucinska, A. (2015) The Art market in the European Union, *International Advances in Economics Research*, 21, 67–79.
- MacKinnon, J. G. and H. L. White (1985) Some heteroskedasticity consistent covariance matrix estimators with improved finite sample properties, *Journal of Econometrics*, 21, 53–70.
- Mandel, B. (2009) Art as an Investment and Conspicuous Consumption Good, *American Economic Review*, 99(4), 1653-63.
- Mei, J. and Moses, M. (2002) Art as an investment and the underperformance of masterpieces, *American Economic Review*, 92, 1656–68.
- Onofri, L. (2009) Old Master Paintings, Export Veto and Price Formation: an Empirical Study, *European Journal of Law and Economics*, 28, 149-161.
- Pesando, J. and Shum, P. (1999) The Returns to Picasso's Prints and to Traditional Financial Assets, 1977 to 1996, *Journal of Cultural Economics*, 23, 183–192.
- Renneboog, L. and Van Houtte, T. (2002) The monetary appreciation of paintings: From realism to Magritte, *Cambridge Journal of Economics*, 26, 331–57.
- Renneboog, L. and Spaenjers, C. (2011) The Iconic Boom in Modern Russian Art, *Journal of Alternative Investments*, 13(3), 67–80.
- Renneboog, L. and Spaenjers, C. (2013) Buying beauty: On prices and returns in the art market, *Management Science*, 59, 36–53
- Silver, M. and Heravi, S. (2007) Why elementary price index number formulas differ: evidence on price dispersion, *Journal of Econometrics*, 140, 874–883.
- Sroufe, R. and Valsan, C. (2006) Hedonic Models and Pre-Auction Estimates: Abstract Art Revisited, *Economics Bulletin*, 26(5), 1–10.
- Stein, J.P. (1977) The Monetary Appreciation of Paintings, *Journal of Political Economy*, 85, 1021-1036.

- Stine, R. (1989) An Introduction to bootstrap methods: examples and ideas, *Sociological Methods & Research*, 18, 243–291, doi:10.1177/0049124189018002003
- Teekens, R. and Koerts, J. (1972) Some Statistical Implications of the Log-Transformation of Multiplicative Models, *Econometrica*, 40(5), 793–819.
- Triplett, J.E. (2004) *Handbook on Hedonic Indexes for Information and Quality Adjustments in Price Indexes*, Working paper 2004/9, OECD Directorate for Science, Technology and Industry, Paris, France.
- Worthington, A. C. and Higgs, H. (2004) Art as an investment: Risk, return and portfolio diversification in major painting markets, *Accounting and Finance*, 44, 257–72.

Table 1 Description of the four data sets and key statistics

	Data set: A	Data set: B	Data set: C	Data set: D
Artist / Group	Surrealists	Impressionists	Pierre-Auguste Renoir	Pablo Picasso
Born–Died	NA	NA	1841–1919	1881–1973
Number of Sales	3,780	2,121	1,972	1,322
Period of Sales	1988-2014	1985-2012	1985–2014	1985-2014
Geometric Mean of Price (US\$)	204,703	894,190	350,045	1,501,646
Geometric Std. Dev. of Price (US\$)	4,484	28,109	10,706	55,697

Table 2 Detailed characteristics and key statistics of the artists included in data sets A and B

Artist	Number of Sales	Born–Died	Geometric Mean Price (US\$)	Geometric Std. Dev. Price (US\$)
Data set A: Surrealists				
Giorgio de Chirico	834	1888-1978	166,815	6,373
Joan Miró	761	1893-1983	385,297	20,281
Max Ernst	628	1891-1976	190,126	8,832
René Magritte	576	1898-1967	578,990	27,905
Roberto Matta	981	1911-2002	84,924	2,582
Data set B: Impressionists				
Alfred Sisley	342	1839–1899	947,317	37,842
Camille Pissarro	586	1839–1903	760,732	37,256
Claude Monet	586	1840–1926	1,946,467	115,671
Odilon Redon	193	1840–1916	177,615	16,528
Paul Gauguin	167	1848–1903	1,190,514	143,954
Paul Signac	247	1863–1935	557,506	52,829

Table 3 Surrealist artists: year-by-year results based on (i) HPM, and, (ii) bootstrap algorithm based on HPM-return estimates after correcting for log-transformation bias

[1] Year	[2] Number of Sales ¹	Estimates from HPM				Estimates from Bootstrap ⁵			
		[3] β	[4] β Std. Error ²	[5] Year-to- year Return ³	[6] Year-to- year Corrected Return ⁴	[7] Year-to- year Corrected Return ⁴	[8] Std. Error Return ⁴	[9] Return ⁴ Lower CL	[10] Return ⁴ Upper CL
1988	110	—	—	—	—				
1989	134	0.501***	0.093	65.06%	71.71%	70.72%	18.88%	35.80%	110.69%
1990	202	0.643***	0.086	15.28%	11.68%	15.47%	15.32%	-11.01%	49.00%
1991	106	0.167	0.090	-37.90%	-44.44%	-50.95%	6.27%	-62.51%	-38.11%
1992	127	0.096	0.089	-6.87%	7.04%	11.29%	13.44%	-12.59%	40.42%
1993	103	-0.253**	0.090	-29.47%	-33.93%	-23.05%	9.37%	-40.43%	-3.90%
1994	127	-0.116	0.091	14.77%	23.88%	18.35%	16.58%	-11.58%	53.59%
1995	110	-0.189	0.104	-7.10%	-1.81%	-8.08%	14.32%	-32.27%	24.02%
1996	125	-0.269**	0.088	-7.69%	-28.14%	-19.66%	11.97%	-41.70%	5.51%
1997	126	-0.364***	0.089	-9.05%	6.16%	3.95%	12.89%	-18.99%	31.43%
1998	144	-0.082	0.094	32.58%	49.42%	50.70%	20.04%	15.04%	93.98%
1999	132	-0.126	0.101	-4.33%	-11.15%	-7.04%	14.43%	-32.52%	25.32%
2000	129	-0.203*	0.089	-7.39%	-25.24%	-26.05%	12.54%	-48.42%	0.12%
2001	128	-0.294**	0.093	-8.68%	12.45%	8.21%	14.43%	-17.21%	37.82%
2002	106	-0.263**	0.098	3.12%	-1.40%	-5.77%	13.74%	-29.07%	25.23%
2003	115	-0.019	0.098	27.65%	32.05%	47.75%	22.03%	6.06%	90.92%
2004	126	0.222*	0.091	27.23%	15.26%	14.14%	17.19%	-14.62%	54.04%
2005	131	0.139	0.084	-7.92%	-8.67%	-8.75%	10.80%	-28.15%	13.54%
2006	136	0.343***	0.090	22.57%	44.13%	47.92%	16.41%	18.87%	82.40%
2007	171	0.446***	0.088	10.84%	4.89%	-0.55%	12.15%	-24.41%	24.57%
2008	176	0.397***	0.089	-4.73%	-7.87%	-5.81%	11.73%	-27.98%	18.46%
2009	183	0.330***	0.084	-6.52%	-12.97%	-13.36%	9.79%	-30.54%	7.67%
2010	166	0.243**	0.086	-8.35%	0.33%	3.99%	11.21%	-16.62%	27.01%
2011	173	0.418***	0.093	19.13%	28.06%	22.06%	17.19%	-7.97%	59.03%
2012	198	0.367***	0.087	-4.95%	-16.47%	-12.07%	12.47%	-35.97%	12.89%
2013	165	0.315**	0.087	-5.10%	-5.92%	-2.03%	11.08%	-21.75%	20.46%
2014	131	0.119	0.088	-17.79%	-17.11%	-13.71%	9.41%	-30.79%	5.22%

¹1988 is the reference year

²Standard error computed using the HCCME HC2 as described in MacKinnon and White (1985)

³Computed directly from the HPM coefficients, before applying the log-transformation bias correction

⁴Return computed after applying the log-transformation bias correction

⁵The bootstrap-algorithm results are based on the HPM return estimates once the bias correction has been applied; CL refers to the confidence interval limits

NOTE: *p<0.05; **p<0.01; ***p<0.0001

Table 4 Surrealist artists: Summary of overall results

Parameter	Estimates from HPM		Estimates from Bootstrap ³		
	Point Estimate ¹	Corrected Point Estimate ²	Point Estimate	Lower CL	Upper CL
Average Annual Return	2.48%	3.54%	4.53%	2.91%	6.31%
Std. Deviation Annual Return	20.89%	26.18%	30.31%	24.51%	37.03%
Risk/Return (RR) ⁴	0.119	0.135	0.148	0.114	0.181
Correlation with S&P 500	0.118	0.155	0.144	0.009	0.286

¹ Computed directly from the HPM coefficients, before applying the log-transformation bias correction

² After applying the log-transformation bias correction

³ The bootstrap-algorithm results are based on the HPM return estimates once the bias correction has been applied; CL refers to the confidence interval limits

⁴ The Risk/Return (RR) metric is the standard deviation of the annual return divided by its mean

Table 5 All four (4) data sets: summary of overall results based on (i) HPM, and, (ii) bootstrap algorithm based on HPM-return estimates after correcting for log-transformation bias

Data Set	Type of Metric	Estimates from HPM					Estimates from Bootstrap ³										
		Adj. R ²	Annual Return	Std. Dev. Annual Return	Risk/Return (RR)	Corr. w/ S&P 500	Annual Return	Difference Corrected HPM vs. Bootstrap ⁴	Annual Return Lower CL	Annual Return Upper CL	Std. Dev. Annual Return	Risk/Return (RR)	RR Lower CL	RR Upper CL	Corr. w/ S&P 500	Corr. Lower CL	Corr. Upper CL
Surrealist artists	Regular ¹ computation	0.718	2.48%	20.89%	0.119	0.118											
	Corrected ²		3.54%	26.18%	0.135	0.155	4.53%	-21.77%	2.91%	6.31%	30.31%	0.148	0.114	0.181	0.144	0.009	0.286
Impressionist artists	Regular ¹ computation	0.718	9.41%	30.46%	0.309	0.435											
	Corrected ²		9.76%	31.64%	0.308	0.210	12.44%	-21.59%	9.39%	16.71%	41.74%	0.299	0.255	0.342	0.184	-0.099	0.462
Pierre-Auguste Renoir	Regular ¹ computation	0.802	5.14%	24.43%	0.210	0.135											
	Corrected ²		5.74%	26.60%	0.216	0.288	7.92%	-27.52%	5.90%	10.30%	35.07%	0.226	0.194	0.260	0.208	0.005	0.405
Pablo Picasso	Regular ¹ computation	0.701	11.47%	34.12%	0.336	0.395											
	Corrected ²		15.57%	43.97%	0.354	0.224	19.52%	-20.20%	13.72%	28.84%	55.57%	0.351	0.307	0.395	0.149	-0.044	0.368

¹Estimates based on the HPM coefficients before applying the log-transformation bias correction

²Estimates based on the HPM coefficients after applying the log-transformation bias correction

³The bootstrap algorithm results are based on the HPM return estimates once the bias correction has been applied; CL refers to the confidence interval limits

⁴Percentage difference between (i) the bootstrap-based point-estimate of annual return and (ii) the HPM-based point-estimate of the annual return after correcting for the log-transformation bias