Duisenberg school of finance - Tinbergen Institute Discussion Paper

TI 13-152/IV/ DSF 61

Does it pay to invest in Art? A Selection-corrected Returns Perspective

Arthur Korteweg[1]
Roman Kräussl[2]
Patrick Verwijmeren[3]

[1] *Stanford Graduate School of Business;*
[2] *Luxembourg School of Finance and the Center for Alternative Investments at Goizueta Business School, Emory University;*
[3] *Erasmus School of Economics, Erasmus University Rotterdam, Duisenberg School of Finance and Tinbergen Institute, The Netherlands; University of Melbourne; University of Glasgow.*

Tinbergen Institute is the graduate school and research institute in economics of Erasmus University Rotterdam, the University of Amsterdam and VU University Amsterdam.

More TI discussion papers can be downloaded at http://www.tinbergen.nl

Tinbergen Institute has two locations:

Tinbergen Institute Amsterdam
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 525 1600

Tinbergen Institute Rotterdam
Burg. Oudlaan 50
3062 PA Rotterdam
The Netherlands
Tel.: +31(0)10 408 8900
Fax: +31(0)10 408 9031

Duisenberg school of finance is a collaboration of the Dutch financial sector and universities, with the ambition to support innovative research and offer top quality academic education in core areas of finance.

DSF research papers can be downloaded at: http://www.dsf.nl/

Duisenberg school of finance
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 525 8579

Does it Pay to Invest in Art?
A Selection-corrected Returns Perspective

Arthur Korteweg, Roman Kräussl, and Patrick Verwijmeren*

Abstract

This paper shows the importance of correcting for sample selection when investing in illiquid assets with endogenous trading. Using a large sample of 20,538 paintings that were sold repeatedly at auction between 1972 and 2010, we find that paintings with higher price appreciation are more likely to trade. This strongly biases estimates of returns. The selection-corrected average annual index return is 7 percent, down from 11 percent for traditional uncorrected repeat-sales regressions, and Sharpe Ratios drop from 0.4 to 0.1. From a pure financial perspective, passive index investing in paintings is not a viable investment strategy, once selection bias is accounted for. Our results have important implications for other illiquid asset classes that trade endogenously.

JEL classification: D44, G1, Z11

Keywords: Art investing; Selection bias; Asset allocation

First Draft: December 13, 2012

Current Draft: July 3, 2013

* Arthur Korteweg (korteweg@stanford.edu) is from Stanford Graduate School of Business, Roman Kräussl (r.g.w.kraussl@vu.nl) is from Luxembourg School of Finance and the Center for Alternative Investments at Goizueta Business School, Emory University, and Patrick Verwijmeren (verwijmeren@ese.eur.nl) is from Erasmus University Rotterdam, University of Melbourne, University of Glasgow, Duisenberg School of Finance and the Tinbergen Institute.
We thank Les Coleman, Elroy Dimson, Bryan Lim, George Parker, and Christophe Spaenjers for useful comments.

For the last two decades, investors are allocating increasingly larger shares of their portfolios to alternative assets. Many of the alternative asset classes, such as private equity and real estate, and even certain traditional assets such as corporate bonds, are highly illiquid. This complicates return evaluation, especially when trades are endogenously related to the performance of the asset. Endogenous trading gives rise to a sample selection problem. This paper demonstrates the empirical first-order importance of correcting for sample selection when evaluating performance and constructing optimal portfolios that include alternative assets.

Among the alternative assets, paintings (and other collectibles) are often considered a comparatively safe investment in times of financial turmoil (Dimson and Spaenjers, 2011). Finding low or negative correlations between art[1] and public equity markets, a growing body of academic research argues that art investments should be included in optimal asset allocations (see, for example, Mei and Moses, 2002; Taylor and Coleman, 2011). Indeed, investors allocate about 6 percent of their total wealth to so-called passion investments[2], and several art funds have been created to allow for diversified investments in art.

The returns and risks of art investments are, however, not beyond debate, and indeed not well known. Constructing an art index and computing the return to art investing is a non-trivial exercise, as prices are not observed at fixed intervals, but only when the artwork trades. Goetzmann (1993, 1996) argues that these trades are endogenous, and he conjectures that paintings that have appreciated in value are more likely to come to market, resulting in high observed returns for paintings that sell, relative to the population. As a result, the observed price appreciation is not representative of the entire market for paintings. In fact, in periods with few

[1] Following the literature, we use the "art" and "paintings" interchangeably throughout the paper.
[2] See the article "Follow your heart" in the Wall Street Journal on September 20, 2010. Passion investments include art, wine, and jewelry.

trades, it is possible to observe high and positive returns even though overall values of paintings are declining.

We present an econometric model of art indices, based on the framework developed by Korteweg and Sorensen (2010, 2012), that generalizes the standard repeat-sales regression (RSR; see Bailey et al. (1963) and Case and Shiller (1987)) to correct for selection bias in the sample of observed sales. This model explicitly specifies the entire path of unobserved valuation and returns between sales, as well as the probability of observing a trade at each point in time, and estimates the selection-corrected price for each individual artwork at each point in time, even when it is not traded. We estimate the model using a unique proprietary auction database from which we construct one of the largest samples of repeat-sales of paintings in the literature to date, with 20,538 paintings being sold a total of 42,548 times between 1972 and 2010.

We find that selection bias is of first-order importance. Paintings are indeed more likely to sell when they have risen in value, consistent with Goetzmann's hypothesis. The difference between our selection-corrected index and the standard (non-corrected) RSR index is economically and statistically large, and robust across specifications. Normalizing indices at 100 in 1972, the RSR index is around 4,300 in 2010, the end of our sample period, whereas the selection-corrected index ends around 1,000 to 1,300, depending on the specification of the selection model. This finding is in line with findings in the real estate markets, where Case et al. (1997) and Korteweg and Sorensen (2012) show that appreciating properties trade faster.

The selection correction has important implications for asset allocation decisions. The annual return to the standard RSR art index over the period 1972 to 2010 is on average 11 percent, with volatility of 13 percent, a Sharpe Ratio of 0.4, and a negative correlation with equity returns. Given these statistics, a mean-variance investor would allocate twice as much of her portfolio to art compared to public equities. She would appear to earn a portfolio Sharpe Ratio of 0.5, which is considerably higher than the Sharpe Ratio of 0.3 on a portfolio of equities only.

However, these returns are not what a passive art investor earns, unless she is able to pick a portfolio of “winners” that rise in value in similar fashion to the paintings that come to auction (such that the RSR index is representative for her portfolio). Instead, it is the selection-corrected returns that are more representative for the experience of an investor who has invested in a well-diversified, passive portfolio of paintings. After correcting for selection bias, we find an average annual art index return of 7 percent, four percentage points lower than the non-corrected index. The selection-corrected Sharpe Ratio is down to 0.1, and the correlation with stock market returns is virtually zero. This makes art a less attractive investment, and the resulting optimal portfolio weight on art is less than half the weight on stocks, with no meaningful improvement in portfolio Sharpe Ratios relative to a portfolio of only stocks.

Our model also allows us to separately assess various styles of paintings. We distinguish between five styles: Post-war and Contemporary, Impressionist and Modern, Old Masters, American, and 19th Century European. In addition, we consider returns to top selling artists. Although we observe interesting differences between the selection-corrected returns per style, and positive portfolio allocations to either Post-war/Contemporary paintings or paintings from top artists, no single style receives non-zero allocations that are robust across all specifications.

To our knowledge, this is the first paper to show the statistical and economic importance of selection effects for price indices of collectibles, and the first to show the importance of selection bias for performance evaluation and portfolio allocation for any asset class. A number of papers in the literature estimate art returns and assess their usefulness for optimal portfolio construction without accounting for selection bias (see Ashenfelter and Graddy (2003) for an excellent overview of this literature, including estimates of art returns from both RSR and hedonic indices in their Table 1). Mei and Moses (2002) and more recently Taylor and Coleman (2011) find that art has attractive characteristics for diversifying investment portfolios, while Renneboog and Spaenjers (2013) estimate average returns that are somewhat lower than the prior literature.

We should note that this paper takes a passive investor view of art index investing. As in real estate, for example, most individuals cannot afford a diversified exposure to the art market through the purchase of individual paintings, and investing in an index is a feasible alternative. With this view, we ignore any aesthetic return on art, i.e., we assume no consumption utility of owning art, since the investor does not have access to the artworks underlying the index. In addition, in order to focus on the effect of selection bias on portfolio weights, we ignore certain frictions in the art market such as transaction costs, insurance costs, and illiquidity costs (see Ang et al., 2013). These frictions all make art less attractive and thus reinforce our main result that there is little benefit from passively investing in paintings. Our result is also robust to considering higher order moments of art index returns, such as skewness and kurtosis.

Our paper focuses on repeat-sales regressions, but the sample selection problem is also important for other price indices that are estimated from sales or auction data, such as the hedonic price index: if artworks that are sold are not representative for the underlying population, then shadow prices estimated from those sales will be biased upwards from their true values in the population. The impact of selection bias on returns in hedonic regressions is an important issue for future work.

Finally, the results in this paper suggest that sample selection in performance evaluation and portfolio allocation is of first-order importance in illiquid markets with endogenous trading, and exploring the strength of this effect in other asset classes is an important avenue for future research.

The remainder of the paper is organized as follows. In Section 1, we describe our model for prices and trades of artworks. Section 2 presents a description of our data. We discuss the art indices in Section 3, both with and without correcting for sample selection, and perform additional tests on the existence and strength of the selection problem in Section 4. In Section 5 we analyze optimal asset allocations, and Section 6 concludes.

1. A selection model of prices and sales of artworks

1.1. Selection model

We decompose the log return of an artwork, *i*, from time *t-1* to *t*, into two components,

$$r_i(t) = \delta(t) + \varepsilon_i(t). \tag{1}$$

The first return component, $\delta(t)$, is the log price change of the aggregate art market from time *t-1* to *t*. We show below how to use the time series of δ to construct a price index for the market. The second return component, $\varepsilon_i(t)$, is an idiosyncratic return that is particular to the individual artwork. We assume that ε has a normal distribution with mean zero and variance σ^2, and is independent over time and across artworks.

If we observe sales of artworks at both time *t-1* and *t*, then the returns are observed, and estimation is straightforward. However, art is sold very infrequently. Typically, years pass between consecutive sales. With a sale at time *t-h* and at time *t*, the observed *h*-period log return is derived from the single-period returns in equation (1) by summation:

$$r_i^h(t) = \sum_{\tau=t-h+1}^{t} \delta(\tau) + \varepsilon_i^h(t). \tag{1}$$

The error term $\varepsilon_i^h(t)$ is normally distributed with mean zero and variance $h\sigma^2$. By defining indicator variables for the periods between sales, the δ's can be estimated by standard GLS regression techniques, scaling return observations by $1/\sqrt{h}$ to correct for heteroskedasticity. This is the repeat-sales regression (RSR) technique that is standard in the literature (Bailey et al., 1963, and Case and Shiller, 1987).

The δ estimates are consistent as long as the indicator variables in the RSR regression are uncorrelated with the error term, i.e., if the probability of a sale is unrelated to the idiosyncratic return component. However, in their survey of the literature, Ashenfelter and Graddy (2003) highlight the concern that art prices may be exacerbated during booms as "better" paintings may come up for sale. Similarly, Goetzmann (1993, 1996) argues that selection biases are important in art data because the decision by an owner to sell a work of art may be conditional on whether or not the value of the artwork has increased.

To correct the repeat-sales model for selection bias of this nature, we specify the sales behavior of art following the model of Korteweg and Sorensen (2010, 2012) that was developed for Venture Capital and real estate. Suppose a sale of artwork *i* at time *t* occurs whenever the latent variable $w_i(t)$ is greater than zero, and remains untraded otherwise, i.e.,

$$w_i(t) = W_i(t)'\alpha_0 + r_i^0(t)\alpha_r + \eta_i(t), \tag{2}$$

where $r_i^0(t)$ is the return since the last sale of the artwork, and is mostly unobserved, except when $w_i(t) > 0$. The vector $W_i(t)$ contains observed covariates. The error term, $\eta_i(t)$, is i.i.d. normal with mean zero and variance normalized to one, and independent of $\varepsilon_i(t)$. The normalization is necessary, but without loss of generality, because the parameters in (3) are only identified up to scale, as in a standard binary probit model.

The selection model, consisting of the observation equation, (1), and the selection equation, (3), nests the classic RSR model. If the selection coefficient, α_r, equals zero, then trades occur for reasons unrelated to price, there is no selection bias, and we recover the standard RSR model. Conversely, if Goetzmann's conjecture is correct and artworks that have risen in value are more likely to sell, then we should find a positive selection coefficient. By estimating and testing the selection coefficient, we allow the data to speak to the importance of selection bias.

From an econometric perspective, the model is a dynamic extension of Heckman's (1979, 1990) selection model. As in Heckman's model, our model adjusts not only for selection on observable variables, such as the size or style of a painting, but also controls for selection on unobservable variables. However, Heckman's model assumes that observations are independent, implying that observations for which price data are missing, are only informative for estimating the selection model, (3), in the first stage, but do not carry any further information for the price index in equation (1) in the second stage. Since prices are path-dependent, this independence assumption fails to hold. Each observation carries information about not only the current price, but also about past and future prices of a painting, even at times when the artwork does not trade. Unlike the standard selection model, our model does not impose the independence assumption, and uses all information to make inference about the price path of individual art works, and the parameters of interest, α, δ, and σ^2.

The downside of allowing for the dependencies between observed and missing data is that it makes estimation more difficult relative to the standard selection model. We use Markov chain Monte Carlo (MCMC), a Bayesian estimation technique, to estimate the model parameters (see Korteweg and Sorensen, 2012, for details).

Our model extends Korteweg and Sorensen (2012) by allowing for separate indices for different styles of artwork, or whether the artist belongs to the top 100 artists. For these specifications, we replace equation (1) with

$$r_i(t) = X_i' \cdot \delta(t) + \varepsilon_i(t). \tag{3}$$

The vector X_i is a set of dummy variables that indicate to which category the painting belongs. The categories need not be mutually exclusive. For example, a painting can belong both to the "Old Masters" style and be in the "Top 100 Artists" category (to be defined below).

1.2. Price indices and portfolio returns

Following the art literature, we focus on the arithmetic price index, $\Pi(t)$, across *N* artworks relative to base year 0,

$$\Pi(t) \equiv \sum_{i=1}^{N} P_i(t) \ / \sum_{i=1}^{N} P_i(0), \tag{4}$$

where $P_i(t)$ is the price of painting *i* at time *t*. The index can be constructed from the selection model estimates by correcting for the Jensen inequality term due to the log operator (see, for example, Goetzmann (1992) and Goetzmann and Peng (2002)),

$$\Pi(t) = \Pi(t-1) \cdot \exp\left(\delta(t) + \frac{1}{2}\sigma^2\right). \tag{5}$$

The return on this index, $\Pi(t)/\Pi(t-1)$, captures the experience of an investor who invests in the aggregate market, or who owns a portfolio that is representative of the aggregate market. It also represents the expected return to an investor who invests a fixed amount in each painting.

2. Art data

We construct a sample of repeat sales from the Blouin Art Sales Index (BASI), an online database that provides data on artworks that are sold at auction at over 350 auction houses worldwide. [3] The BASI database is presently the largest known database of artworks, containing roughly 4.6 million works of art by more than 225,000 individual artists over the period 1922 to 2010. We solely focus on paintings, which represent 2.3 million artworks in the database.

For each auction record, the database contains information on the artist, the artwork, and the sale. We observe the artist's name, nationality, year of birth, and year of death (if applicable). For the artwork, we know its title, year of creation, medium, size, and style, and whether it is signed or stamped. For the sale, we have data on the auction house, date of the auction, lot number, hammer price (the price for which the artwork was sold, converted to U.S. dollars at the prevailing spot price), and whether the artwork has been bought in or was withdrawn.[4]

We identify repeat sales by matching auction sales records using artists' names, artwork names, painting size, and medium (similar to the matching procedures in Taylor and Coleman, 2011, and Spaenjers and Renneboog, 2013). We start our search in 1972, due to the limited coverage of the database before that time. We delete buy-ins to avoid Goetzmann's (1993) concern that particular auction records are wrongly classified as sales when the painting fails to meet the seller's reservation price. To eliminate false matches, we remove paintings from the same artist with the generic titles "untitled" and "landscape." We further check whether the remaining potential repeat-sales are true repeat-sales by manually searching for the artwork's provenance, which shows the chronology of the artwork's earlier sales. The provenance is typically found on the websites of the auction houses. For instance, Christie's and Sotheby's

[3] Art is not only sold in auction but also privately, for example through dealers. Renneboog and Spaenjers (2013) note that it is generally accepted that auction prices set a benchmark that is also used in the private market.
[4] An artwork is "bought in" when the bidding does not reach the reserve price, and the artwork goes unsold.

provide online provenance information on all auction sales since 1998. When we are in doubt about whether we are dealing with a true repeat sale, we delete it from our sample. Our final sample includes 42,548 sales of 20,538 unique paintings.

Figure 1 Panel A shows the number of sales in the repeat-sales sample, broken down by the number of first, second, and third or more sales for the artworks in our repeat-sales sample.[5] For comparison, Panel B shows the number for the full BASI dataset since 1972. The full sample shows substantial growth in sales over time, peaking in 2006, whereas the repeat-sales sample has the highest number of sales in 1989. This difference is due to the drop in the number of first sales, as paintings that sell for the first time in the later part of our sample period have a smaller probability of being sold for a second time by the end of the sample period and thus have a smaller chance of being included in the repeat-sales sample.

[Please insert Figure 1 here]

Panel A of Table 1 shows descriptive statistics of the paintings in our repeat-sales sample. The average hammer price in the full sample is $61,939, with a long right tail of extremely expensive paintings. The average surface of the paintings is about 547,000mm^2, or 0.55m^2. Around 22 percent of sales take place at Christie's auction house, and 25 percent at Sotheby's. For 20 percent of sales, the auction house is located in London, and another 20 percent are sold in New York. Using the same style classifications as Christie's and Sotheby's, the BASI database distinguishes between six broad styles. The Impressionist and Modern style accounts for one third of sales, followed by European 19th Century paintings with one fourth of sales. About 16 percent of sales are of the Post-war and Contemporary style, 12 percent are American paintings, and 5 percent are Old Masters. The residual "Other style" category makes up the

[5] Observing more than three sales is extremely rare in our sample period: we only observe 26 paintings with four sales, and one painting with five sales.

remaining 9 percent of sales. Nearly 10 percent of sold paintings are by artists with total dollar sales in the top 100 of the BASI database over the sample period.[6] Finally, more than two percent of sales occur within two years after the artist has deceased.

[Please insert Table 1 here]

Panel A also shows the descriptive statistics for the full BASI sample over the same period. Compared to the full sample, more expensive paintings of higher quality are more likely to be sold repeatedly, underscoring the importance of correcting for sample selection. It should be noted that even if the repeat-sales sample were statistically indistinguishable from the full sample of sales, the sample selection issue that we address in this paper may still be present, as even the full sample of sales may not be representative of the underlying population of paintings.

Panel B provides information about the sale-to-sale returns in the repeat-sales sample. The arithmetic price increase between two consecutive sales of the same painting is 123.5 percent on average. The median return is 42.4 percent, and the standard deviation is 368.5 percent. With an average time between sales of 7.6 years, this translates to an average (median) annualized return of 16.5 percent (7.5 percent), with a standard deviation of 32.7 percent. Log returns are lower, on average 43.9 percent (6.9 percent annualized), with a median of 35.3 percent (5.7 percent annualized) and a standard deviation of 78.1 percent (16.7 percent annualized).

[Please insert Figure 2 here]

Figure 2 shows the distribution of the annualized sale-to-sale returns. Although the average return is positive, the distribution shows negative returns occur regularly. Annualized returns below –30 percent or above 70 percent are rare.

[6] We keep the artists in the top 100 category fixed throughout the sample period.

3. Art indices

In this section we first present the results of estimating art indices without taking selection into account. Then we turn to our selection-corrected indices.

3.1. Art indices without selection correction

Figure 3 plots two estimated arithmetic art price indices that do not take selection into account. The first is a standard repeat-sales regression estimated by GLS, weighing each observation by the inverse of the square root of the time between trades, to correct for potential heteroskedasticity as described above (the "GLS index"). The second is a MCMC specification that ignores selection by forcing α_r in equation (3) to equal zero (the "MCMC index"). We assign an index value of 100 to the year 1972 and construct annual end-of-year arithmetic indices as shown in equations (5) and (6).

[Please insert Figure 3 here]

The GLS and MCMC indices practically coincide, mitigating concerns about distributional assumptions of the MCMC estimator. Over the early part of the sample period, the indices rise until they peak at 1,300 in 1990. After bottoming out at 900 in 1993, following the Japanese real estate crisis in the late 1980s/early 1990s, the indices climb again until peaking at around 4,300 in 2007, showing particularly high growth after 2001. In 2010, the end of our sample period, the price indices have largely recovered from the dip in the global financial crisis of 2008/09, and are nearing the 4,300 level again.

3.2. Selection models

The selection models require us to take a stance on what drives the sale of an artwork. We estimate three specifications of the selection equation. All models include the log return since last sale, to capture the direction and strength of the sample selection effect. We also include the time since the last sale, both linearly and squared, in all specifications. Time since last sale functions as an instrument to identify the model from more than distributional assumptions alone: it changes the probability of a sale in the next period without affecting the return on the artwork going forward, based on the common-place assumption that prices incorporate all available public information, which includes the date that the painting last sold. Other variables that may be important for the probability of sale are the size of the painting, whether the artist deceased in the past two years, and the growth in worldwide GDP. The size of the painting may be related to the probability of sale, because smaller paintings are easier to hang and transport. The death of the artist might be relevant as there is a popular belief that artworks are more likely to be sold when the artist has recently deceased. We include worldwide GDP growth as Goetzmann (1993) and Goetzmann et al. (2011) establish an important relation between art and wealth.

[Please insert Table 2 here]

Table 2 shows the estimated coefficients of the selection models. Model A only includes the log return and the time variables. Most importantly, paintings with higher returns since the prior sale are more likely to sell, and hence appear more frequently in the sales data, confirming that selection is important in evaluating art returns. As a result, standard indices exaggerate the price appreciation of the overall market.

The magnitude of the selection coefficient is not only statistically significant, it is also economically meaningful. For example, a painting that was last sold one year ago and has not changed in value since, has a 3.6 percent probability of being sold in the next year. Had the painting increased in value by one standard deviation of the annualized return, or 16.7 percent, then the probability of sale would be 4.1 percent. For a two standard deviation increase in price, the sale probability is 4.7 percent.

Controlling for price appreciation, the time since the prior sale has a non-linear relation to the probability of a sale. An artwork that has traded very recently is less likely to sell again, but as the time since last sale increases, the probability of a sale rises as the coefficient on the squared time since sale dominates. All time effects are significant at the one percent significance level, suggesting that the time variables improve the statistical identification of the model.

The estimates for Model B show that our main finding that returns and the probability of sale are positively related, is robust to adding additional variables to the selection model. The coefficients on these variables provide further insight into what drives sales behavior. A painting is more likely to sell within two years after the artist deceases, and when world GDP is declining. The latter result is a bit difficult to interpret given that we also control for price appreciation separately, but it is in line with situations in which owners are forced to sell in bad times (Campbell, 2008). We do not find a significant effect of a painting's size.

Model C allows for different selection coefficients and intercepts by decade, where we group the 1970s with the 1980s (due to data sparseness), and we split the 2000s in the pre and post financial crisis years. Thus, our periods are 1972 to 1989, 1990 to 1999, 2000 to 2006, and 2007 to 2010. We find evidence for selection in every sub-period. The effect of returns on the probability of trade is especially large in the two periods after 2000, indicating that the selection effect has become more important since the turn of century.

We use Models A, B, and C to construct three selection-corrected indices, which we denote Indices A, B, and C, respectively. Panel A of Figure 4 shows the selection-corrected price indices over time. The selection correction is quite robust across models, as the differences among the three selection-corrected indices are not very large. Panel A also includes the MCMC model without selection-correction. The differences between the non-corrected and the selection-corrected indices are striking. First, the selection-corrected indices are considerably lower than the non-corrected indices, conform the intuition about the effect of a positive selection coefficient. The peak in 1990, which occurred at a 1,300 index level in the non-selection corrected model, occurs at around 850 in the corrected indices. The 2007 peak is around 1,300 (or 1,000 for model C), rather than the 4,300 of the non-corrected model. Second, the selection-corrected indices show an additional peak around 2003, which does not occur in the non-corrected indices. Third, the selection-corrected indices do not recover after the global financial crisis of 2008, unlike the non-corrected indices.

[Please insert Figure 4 here]

In Panel B, we plot the difference between the natural logarithm of the non-selection corrected index and each of the selection-corrected indices. The graph shows that the deviation between the non-selection and the selection-corrected indices already starts in the first years of our sample period, increases steadily over time, and accelerates in the new millennium.

Next, we consider potential differences between art styles. Buelens and Ginsburgh (1993) find differential performance among styles, and possibly, different styles are in favor in different periods. Figure 5 plots the estimated selection-corrected indices.

[Please insert Figure 5 here]

We observe that the price index of Post-war and Contemporary paintings peaks around 1990 and 2007. Impressionist and Modern paintings show large increases in the early period but are hit heavily in 1990, which is in line with the popular interpretation of art observers that the Japanese real estate bubble (which burst in 1990) and the corresponding strong yen in the 1980s had strong effects on the prices of Impressionist paintings (see for example Wood, 1992). Old Masters do not increase as much in value as the other styles over the sample period.

Figure 5 also shows the selection-corrected price index for paintings of the top 100 artists in terms of the total value of sales in our sample period. Top artists outperform all of the styles, and the index peaks both around 1990 and 2007. This result relates to the "masterpiece effect", the general belief among art dealers and critics that highly priced paintings are the best buy (e.g., Adam (2008)). Several prior academic studies examine masterpieces (Mei and Moses, 2002; Renneboog and Spaenjers, 2013), but generally find that masterpieces underperform (Pesando, 1993; Mei and Moses, 2002). Our results suggest that selection bias is an important determinant of this discrepancy, as not controlling for selection biases artificially drives up the returns for artworks with which masterpieces are compared. Paintings from top artists do experience relatively volatile returns, so it is not clear whether they should get higher weights in optimal portfolios. We discuss portfolio allocations after we present further evidence in support of the sample selection problem in the next section.

4. Further supporting evidence for sample selection

The results from the econometric model show that paintings are more likely to trade when they have experienced a high return since prior sale, and that correcting for the sample selection problem is economically important for index construction. In this section we show further corroborating evidence of the existence and strength of the selection problem in the raw data, without relying on the econometric machinery of the selection model.

First, we consider the relationship between the time between sales and the annualized sale-to-sale return. For each painting in the repeat-sales sample, we compute the annualized log return between two adjacent sales, and the number of years between the two sales. Both these quantities are observed in the raw data. In Figure 6, we graph the average annualized sale-to-sale return against the return horizon. For example, for all sale-to-sale returns that occurred over a span of one year or less, the average annualized return is 13 percent, compared to 8 percent for sales that took place between one and two years apart. For longer return horizons the average annualized return is even lower. We also show the median annualized return against the return horizon, which follows a nearly identical pattern.

[Please insert Figure 6 here]

If there were no sample selection problem (i.e., if the selection coefficient, α_r, equals zero in the econometric model), then there should be no systematic relation between annualized returns and the time between sales, and Figure 6 should show a flat, horizontal line. Instead, the line is downward-sloping, which is consistent with a sample selection problem in which paintings with high returns are more likely to trade. To take an oversimplified example, suppose that paintings trade as soon as the return since last sale hits a fixed threshold, say, 10 percent (not annualized). Then the paintings that happen (by chance) to have a high return soon after the prior sale, will trade quickly and show a high annualized return. The paintings that are slower to hit the threshold will trade at a later date and exhibit lower annualized returns.

Figure 6 also shows that the selection problem is plausibly large. The mean annual index return from RSR regressions is essentially an average of the annualized observed returns over all horizons. In a selected sample where the paintings with higher returns are more likely to sell, the selection-corrected average annualized return must be lower than the observed returns, and thus lower than even the long-horizon returns. Based on the Figure 6, a difference of a few percentage points in the mean annual return between the RSR and the selection-corrected indices is not surprising.

The second piece of evidence suggesting sample selection is the relation between annualized returns and trading intensity. In Figure 7, we plot the time series of annualized log returns, where for a given year we average over all sale-to-sale returns for which the second sale is in that year. We only show the time series starting in 1980, because for many sales in the 1970s the first sale takes place prior to the start of our sample in 1972, and we see only the short-horizon returns in the 1970s. In the same graph, we show the time series of trading intensity, defined as the percentage of paintings that sold in the calendar year, calculated from the full BASI dataset of 2.3 million observations (the results are very similar if we use the repeat-sales sample instead).

[Please insert Figure 7 here]

Without selection, i.e. when paintings trade for reasons unrelated to returns, there should not be a systematic relation between trading intensity and annualized returns. This is clearly not what we observe in Figure 7, which shows a strong positive correlation between trading intensity and the annualized returns. This relation is consistent with the sample selection problem identified in this paper, where a positive shock to the value of paintings results in more paintings that are likely to trade, and a higher average (annualized) realized return, and vice versa for negative shocks.

Our third exercise further exploits the variation in the data, by considering the relation between trading intensity and annualized returns for individual styles of paintings. Consistent with the aggregate results, Table 3 Panel A shows that both the average and median annualized return of a style of painting are strongly positively correlated with the market share of that style, where market share is defined as the number of paintings of a style that sold over the year relative to the total number of paintings that traded across all styles. Panel B shows that this correlation is highly statistically significant in a pooled regression analysis that controls for style fixed effects.

[Please insert Table 3 here]

To summarize, the evidence in this section is supportive of the result from the econometric model that the probability that a painting trades is positively related to its return since the prior sale.

5. Optimal portfolio allocation

In this section, we show that our estimated indices provide important insights regarding the role of paintings for diversification and optimal portfolio allocation. More broadly, the results underscore the importance of adjusting for sample selection for performance evaluation and portfolio optimization in the presence of illiquid assets when trading is endogenous.

We start our analysis with descriptive statistics of the returns to the art indices. Table 4 reports means, standard deviations and Sharpe Ratios of arithmetic and log annual returns for the art indices. For clarity, it should be noted that these are the returns to the indices, and thus are different from the sale-to-sale returns reported in Table 1 and Figures 2, 6, and 7.

[Please insert Table 4 here]

Panel A shows that the standard repeat-sales GLS index that does not correct for selection has an average arithmetic annual return of 11.1 percent with a standard deviation of 12.8 percent, and an annual Sharpe ratio of 0.39. The non-selection corrected MCMC index returns are nearly identical to the GLS index. In contrast, the selection-corrected indices have considerably lower average arithmetic returns, ranging from 6.7 percent to 7.3 percent depending on the selection model specification. The standard deviations are similar to the non-selected indices, around 12.7 percent to 13.3 percent. Consequently, the selection-corrected Sharpe Ratios are considerably lower than the non-corrected indices, ranging from 0.05 to 0.10. Log returns follow a pattern that is similar to the arithmetic returns, but with lower average returns, standard deviations and Sharpe Ratios.

It is important to point out that, as a measure of performance, the standard non selection-corrected indices implicitly assume that an investor can either pick "winners" that rise in value and are thus more likely to sell, or assume that there is no selection problem and that all other holdings of the investor follow the same price path as the paintings that are auctioned off. The selection-corrected indices do not make such assumptions but rather approximate the rise in value of the overall portfolio of paintings, both those that sold and those that did not. The selection-corrected returns are therefore more representative of the experience of an investor who has invested in a well-diversified, passive portfolio of paintings (as noted by Renneboog and Spaenjers (2013), the volatility of investors' portfolios is likely higher if they are less diversified).

Panel A of Table 4 also shows descriptive statistics for the sample period returns of a broad portfolio of U.S. equities (the CRSP value-weighted index including distributions), and the prevailing 1-year U.S. Treasury bill rate at the beginning of the year. Since we use the Treasury rate as our risk-free asset, we do not report its Sharpe Ratio in Table 4.

Despite the low Sharpe Ratios on the selection-corrected art indices relative to the Sharpe Ratio on stocks of 0.30, investing in paintings may still be useful for constructing optimal portfolios if the correlations with stocks are low. Indeed, Panel B of Table 4 shows that the art indices that do not control for selection are robustly negatively correlated with stock returns. For example, the correlation coefficient between the GLS arithmetic art return and equity is -0.08. For comparison, Taylor and Coleman (2011) find strong negative correlations (of about -0.30) between stocks and aboriginal art. Renneboog and Spaenjers (2013) find a correlation coefficient of -0.03 between art and the S&P 500 index. However, Panel B also shows that once we correct for sample selection, the art indices are virtually uncorrelated with stock returns, with correlations that are less than 0.02 in absolute magnitude.

To examine the portfolio allocation decision more formally, we construct optimal portfolios based on the following assumptions that are common in the literature. First, investors have mean-variance utility, and allocate their portfolio among the risk-free asset, a well-diversified stock index (the CRSP value-weighted market index including distributions), and a well-diversified, passive art index. Second, borrowing and short sales are not allowed. Third, there are no transaction costs to constructing the indices. Fourth, there is no illiquidity return premium on paintings. Fifth, investing in the art index does not provide the investor with access to the artworks underlying the index, and we therefore do not consider any consumption utility of owning art.

[Please insert Table 5 here]

Table 5 Panel A shows the portfolio weights for the tangency portfolio of stocks and art (i.e., the portfolio with the maximum Sharpe Ratio in the presence of a risk-free asset). We focus our discussion on the arithmetic returns. An investor who does not correct for selection bias in art returns assigns considerable weight to art. Based on the arithmetic GLS returns, art receives a portfolio weight of 65 percent, with the remainder assigned to stocks. The portfolio Sharpe ratio of 0.51 is considerably higher than the Sharpe Ratio of 0.30 that is achieved with stocks alone. The MCMC index that ignores selection gives weights that are nearly identical to the GLS index, and for brevity we omit them from the portfolio weights tables. The non-selection corrected indices thus suggest that paintings should play an important role in asset allocation.

In contrast, an investor who corrects for sample selection assigns a significantly lower weight to paintings of 20 percent to 35 percent, depending on the selection model. The portfolio Sharpe Ratio of 0.30 to 0.32 is about 40 percent lower than the Sharpe Ratio of 0.51 from the non-corrected indices, and close to the Sharpe Ratio of the pure stock index, indicating that little is gained by allocating a share of the portfolio to paintings, despite the non-zero portfolio weight.

Panels B, C, and D of Table 5 show the portfolio allocations to paintings, stocks, and the risk-free asset for a mean-variance utility investor with a risk aversion coefficient equal to two, five, and ten, respectively. The results are similar to the tangency portfolio results and underscore our main result: an investor who does not correct for sample selection would allocate nearly twice as much of her portfolio to paintings compared to public equity. With the selection correction, the same investor would put less than half of the weight on paintings relative to stocks, and realize virtually no gain in Sharpe Ratios compared to a portfolio of only public stocks. The Sharpe Ratios are naturally the same across all panels of Table 5, barring minor deviations due to borrowing constraints in Panel B.

The results for log returns are qualitatively similar to the arithmetic returns, albeit with lower weights on paintings for the selection corrected indices.

The remainder of this section shows the robustness of the portfolio allocation result. For the first robustness test, we consider whether investing in particular styles of painting, or in top-selling artists, would be beneficial for forming portfolios, even after controlling for selection.

[Please insert Table 6 here]

Table 6 Panel A reports the results for the different art styles. The only style that receives a non-zero portfolio allocation is the Post-war and Contemporary style. Investing in this style and stocks yields a portfolio Sharpe Ratio of around 0.38. Panel B further includes the index of the Top 100 artists, which drives out the allocation to Post-war and Contemporary paintings in the arithmetic returns, but not in the log returns. Altogether, Table 6 shows that none of the styles or the Top 100 artists receive a robustly positive portfolio weight when controlling for selection.

In our second robustness test, we correct for non-synchronous trading using Dimson's (1979) method with one year leads and lags of stock returns. This may be important as art indices aggregate pricing information over the calendar year, while stock returns are exact year-to-year changes. Goetzmann et al. (2011) show the importance of lagged equity returns on art prices. Table 7 shows that the Dimson-corrected portfolio weights on paintings are lower compared to the weights without the non-synchronicity correction, but the main result is robust: there is little gain in Sharpe Ratios by including art in a portfolio with public equities, after correcting for sample selection.

[Please insert Table 7 here]

In other, non-tabulated robustness tests, we confirm that the main portfolio result is robust to using the equally-weighted CRSP index or the S&P composite index instead of the value-weighted CRSP index, to using the longer 1926 to 2010 period to estimate the average market return and its standard deviation (computing the covariance with art from the correlation over the 1972 to 2010 period but using the standard deviation from the longer time series)[7], to using the return to the Citigroup World Government Bond Index from Datastream instead of the 1-year T-

[7] For the 1926 to 2010 period we need to use the 1-month T-bill rate to compute excess market returns, as the 1-year rate is only available starting 1959.

bill rate, and to adding more asset classes to the optimal portfolio (including real estate returns, hedge fund returns, private equity returns, and commodity returns).

It is important to note that we have abstracted away from certain frictions in the art markets. One important friction that we omitted is transaction costs: the typical buyer's premium in art is up to 17.5 percent of the hammer price, and there are storage and insurance fees. These costs make paintings less appealing for optimal portfolio allocation and hence reinforce our main result. Moreover, Ang et al. (2013) show that illiquidity costs, another potential friction, reduce portfolio allocations, further strengthening our result.

Finally, investors may care about higher moments of returns (for example, if they have power utility rather than mean-variance utility), such as skewness (e.g. Ball et al., 1995, and Harvey and Siddique, 2000) and kurtosis. However, art returns are significantly negatively skewed and have excess kurtosis (results not tabulated), which only makes art less attractive to a power utility investor.

In sum, after correcting for sample selection, we find that paintings play little role in raising the Sharpe ratios of optimal portfolios relative to public equities. The additional frictions in art markets only serve to reinforce this result.

6. Conclusion

We estimate an empirical model that adjusts for selection bias in illiquid asset markets with endogenous trading, using a large dataset of auction sales of paintings. We find a large selection effect of the kind hypothesized by Goetzmann (1993, 1996), namely that paintings that have increased in value are more likely to sell. This has a first-order impact on art indices, lowering the average annual price increase from 11 percent for a standard repeat-sales index to 7 percent for selection-corrected indices, resulting in a drop in annual Sharpe Ratios from 0.4 to 0.1.

Ignoring the sample selection problem would lead investors to allocate more than half their portfolio to art, and they appear to reap large portfolio Sharpe Ratios on the order of 0.5 annually. However, after correcting for selection, the portfolio allocation to art drops dramatically and there is little gain in portfolio Sharpe Ratios relative to the 0.3 Sharpe Ratio of pure public equities. We conclude that investing in a passive art index is unattractive, even without considering transaction and insurance costs, and the risks of forgeries, thefts, and physical damage, unless investors are able to pick winners or there is substantial non-monetary utility from owning and enjoying art.

To our knowledge, this is the first paper to show the importance of the endogenous trading sample selection problem for performance evaluation and optimal portfolio allocation. It stands to reason that other illiquid asset classes exhibit a similar selection problem, and evaluating these other asset classes is an important task for future work.

References

Adam, G., 2008. When Brueghel met Schnabel. Financial Times (21-04-2008).

Ang, A., Papanikolaou, D., and Westerfield, M., 2013. Portfolio choice with illiquid assets. *Working paper*, Columbia University, Northwestern University and University of Washington.

Ashenfelter, O. and Graddy, K., 2003. Auctions and the price of art. *Journal of Economic Literature* 41, 763-786.

Bailey, M.J., Muth, R.F., and Nourse, H.O., 1963, A regression method for real estate price index construction. *Journal of the American Statistical Association* 58, 933–942.

Ball, R., Kothari, S.P., and Shanken, J., 1995, Problems in measuring portfolio performance: An application to contrarian investment strategies. *Journal of Financial Economics* 38, 79-107.

Buelens, N. and Ginsburgh, V., 1993. Revisiting Baumol's "art as a floating crap game". *European Economic Review* 37, 1351-1371.

Case, K.E., Pollakowski, H.O., and Wachter, S.M., 1997. Frequency of transaction and house price modeling. *Journal of Real Estate Finance and Economics* 14, 173-187.

Case, K.E. and Shiller, R.J., 1987, Prices of single family homes since 1970: New indexes for four cities. *New England Economic Review Sept/Oct*, 46–56.

Campbell, R.A.J., 2008. Art Finance. In: Fabozzi, F.J., Handbook of Finance: Financial Markets and Instruments, pp. 605-610, John Wiley & Sons, New Jersey.

Chanel, O., Gérard-Varet, L., and Ginsburgh, V., 1996. The relevance of hedonic price indices: the case of paintings. *Journal of Cultural Economics* 20, 1-24.

Dimson, E., 1979. Risk measurement and infrequent trading. *Journal of Financial Economics* 7, 197-226.

Dimson, E., and Spaenjers, C., 2011. Ex post: The investment performance of collectible stamps. *Journal of Financial Economics* 100, 443-458.

Goetzmann, W., 1992. The accuracy of real estate indices: Repeat sales estimators. *Journal of Real Estate Finance and Economics* 5, 5-53.

Goetzmann, W., 1993. Accounting for taste: Art and financial markets over three centuries. *American Economic Review* 83, 1370-1376.

Goetzmann, W., 1996. How Costly Is the Fall from Fashion? Survivorship Bias in the Painting Market. In *Economics of the Arts: Selected Essays*. Vol. 237, Contributions to Economic Analysis, ed. Victor A. Ginsburgh and Pierre-Michel Menger, 71–84. New York: Elsevier.

Goetzmann, W., Renneboog, L., and Spaenjers, C., 2011. Art and money. *American Economic Review* 101, 222-226.

Goetzmann, W., and L. Peng, 2002. The bias of the RSR estimator and the accuracy of some alternatives. *Real Estate Economics* 30, 13-39.

Harvey, C.R., and A. Siddique, 2000, Conditional skewness in asset pricing tests. *Journal of Finance* 55, 1263-1295.

Heckman, James, 1979, Sample selection bias as a specification error. *Econometrica* 47, 153-162.

Heckman, James, 1990, Varieties of selection bias. *American Economic Review* 80, 313-318.

Korteweg, A., and Sorensen, M., 2010. Risk and return characteristics of venture capital-backed entrepreneurial companies. *Review of Financial Studies* 23, 3738-3772.

Korteweg, A., and Sorensen, M., 2012. Estimating loan-to-value and foreclosure behavior. Working paper.

Markowitz, H., 1952. Portfolio Selection. *Journal of Finance* 7, 77-79.

Mei, J. and Moses, M., 2002. Art as an investment and the underperformance of masterpieces. *American Economic Review* 92, 1656-1668.

Pesando, J.E., 1993. Art as an investment: The market for modern prints. *American Economic Review* 83, 1075-1089.

Renneboog, L., and Spaenjers, C., 2013. Buying beauty: On prices and returns in the art market. *Management Science* 59, 36-53.

Taylor, D., and Coleman, L., 2011. Price determinants of Aboriginal art, and its role as an alternative asset class. *Journal of Banking and Finance* 35, 1519-1529.

Wood, C., 1992. The Bubble Economy. The Japanese Economic Collapse. London: Sidgwick & Jackson.

Figure 1: Number of sales

This figure shows the number of auction sales of paintings in the repeat-sales sample (panel A) and the full sample (panel B) by calendar year. In the repeat-sales sample in panel A, we distinguish between the first, second, and third or more sales of an artwork.

Panel A: Repeat-sales sample

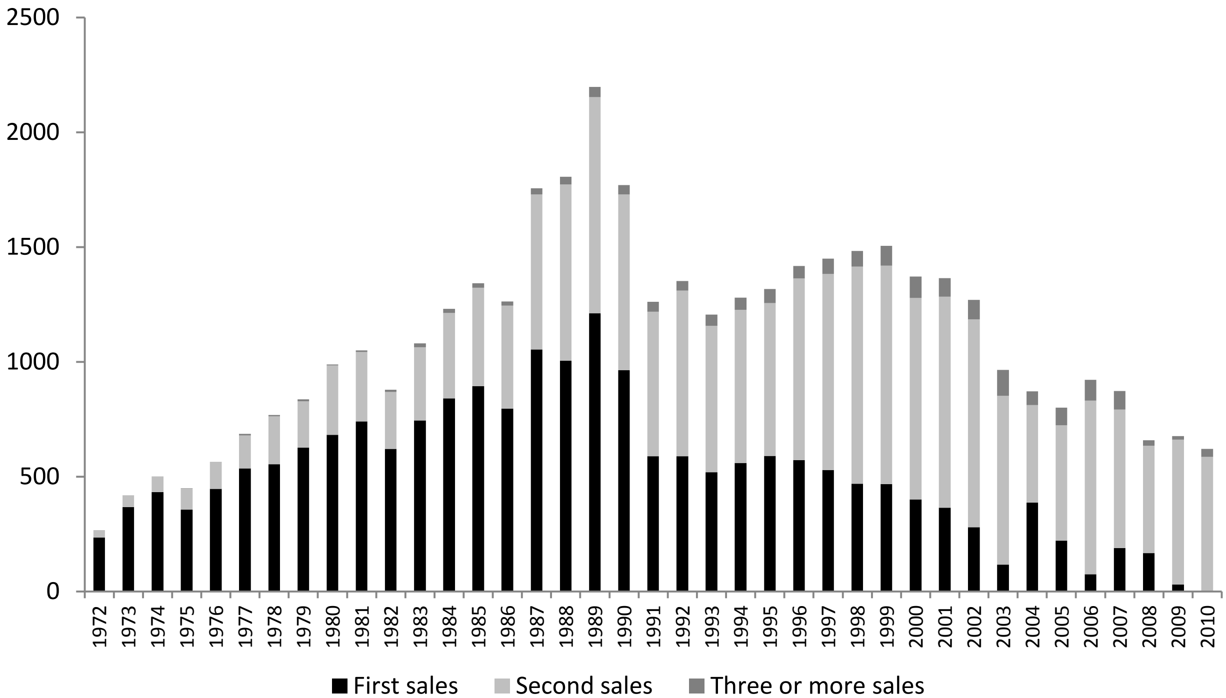

Panel B: Full sample

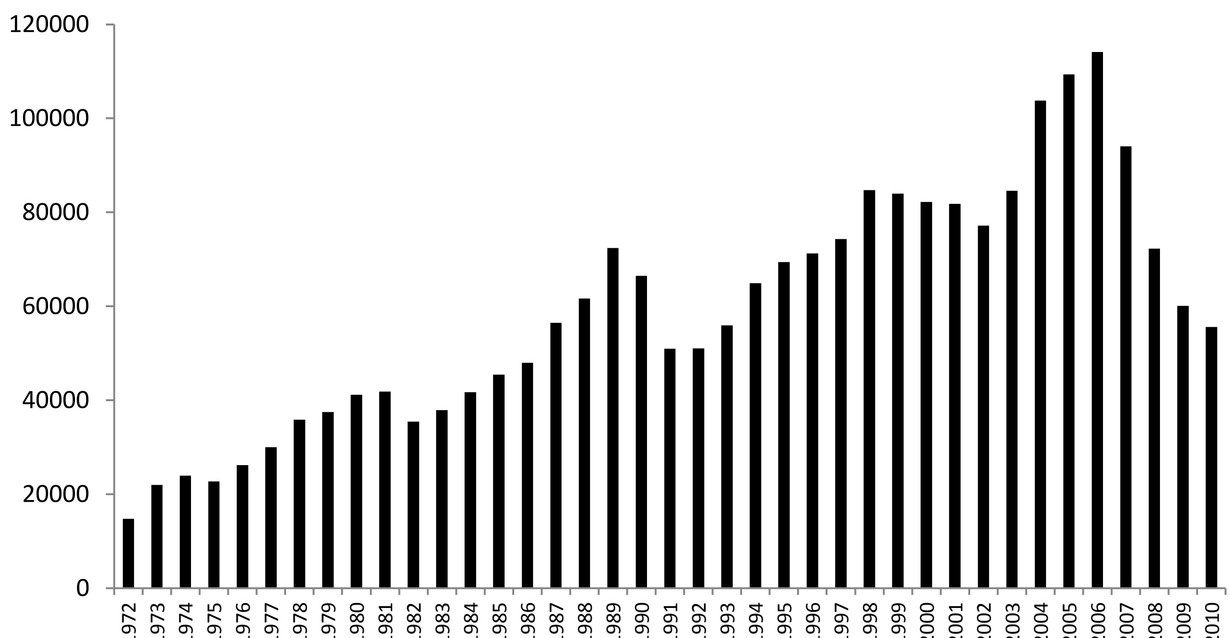

Figure 2: Annualized sale-to-sale return distribution

This figure shows histograms of the annualized sale-to-sale returns of paintings in the repeat-sales sample. Panel A shows the annualized arithmetic sale-to-sale returns, and Panel B shows the natural logarithm of the annualized sale-to-sale returns.

Panel A: Arithmetic returns

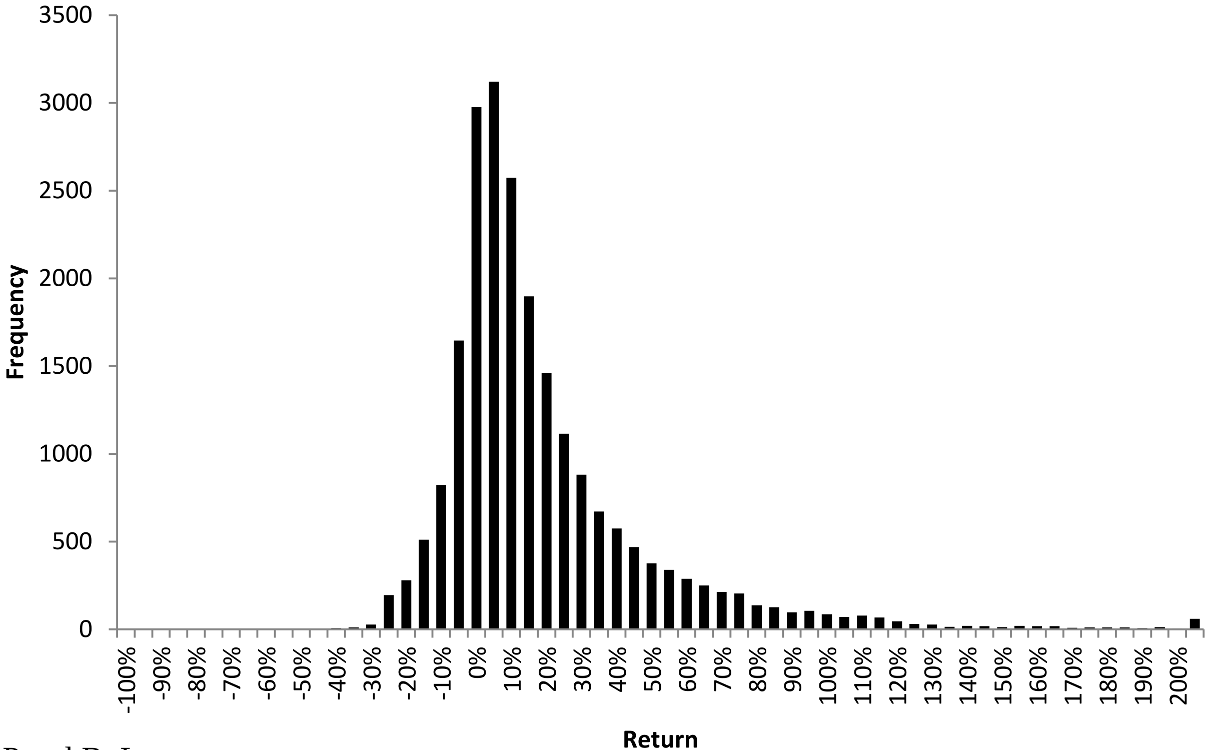

Panel B: Log returns

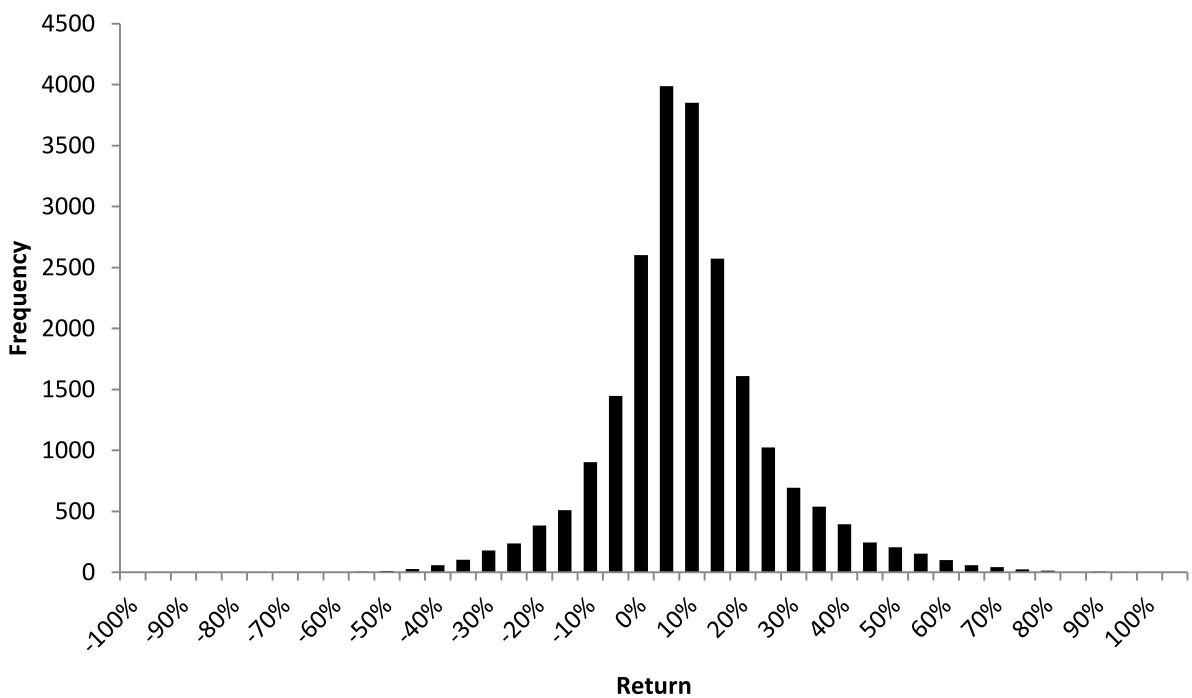

Figure 3: Non-selection corrected price indices

This figure shows repeat-sales arithmetic price indices that do not correct for sample selection. The indices are normalized to an index value of 100 in 1972. The *GLS* index is the standard repeat-sales regression indices as estimated by Generalized Least Squares, with weights that are inverse proportional to the square root of the time between sales. The *MCMC* index is the index estimated by the Markov chain Monte Carlo algorithm when the sample selection problem is forcibly ignored, i.e., α_r in equation (3) is set to zero.

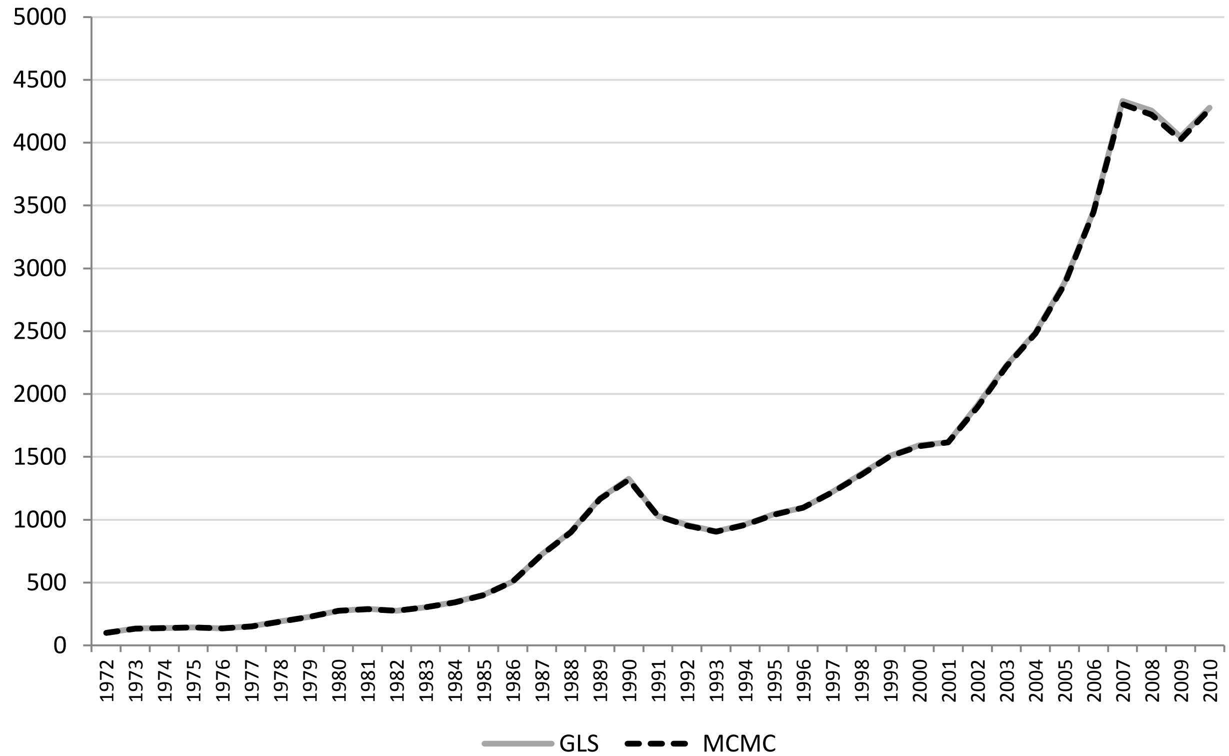

Figure 4: Selection-corrected price indices

Panel A shows price indices corrected for sample selection, and normalized to an index value of 100 in 1972. Models A through C correspond to the specifications of the selection equation as shown in Table 2. For comparison, the figure also shows the MCMC non-selection corrected index over this same period (denoted *No selection*. This is the same index as the *MCMC* index in Figure 3). Panel B graphs the difference between the natural logarithm of the *No selection* index and the selection-corrected indices.

Panel A: Price indices

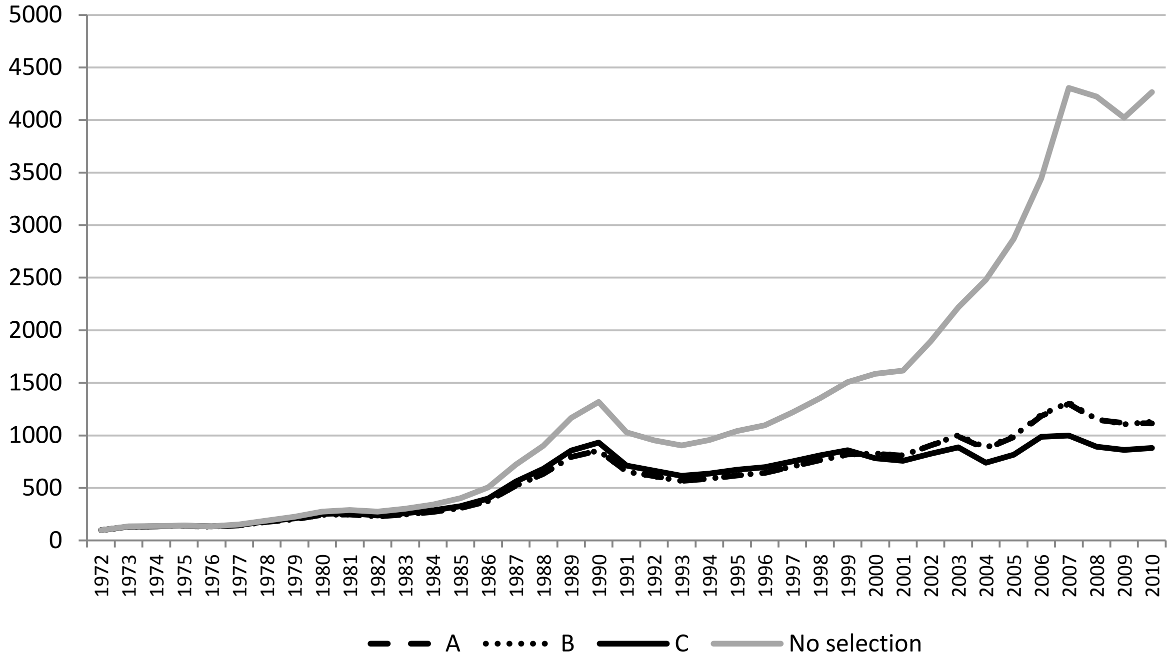

Panel B: Logarithm of No selection index minus logarithm of selection-corrected price indices

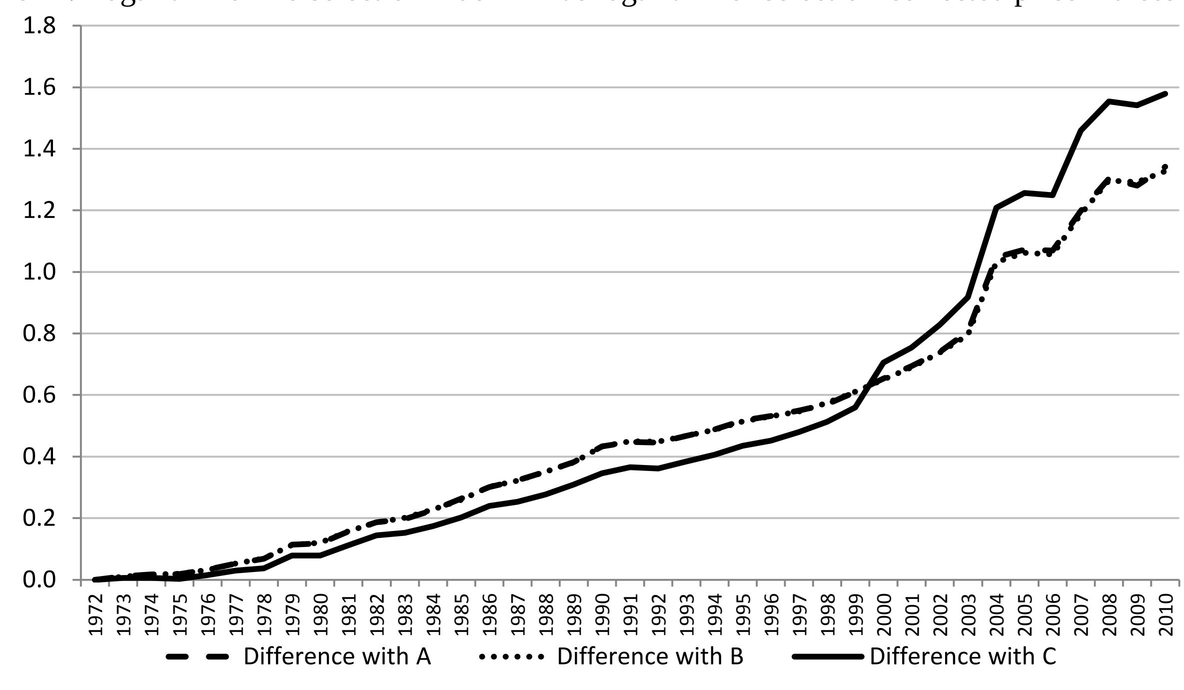

Figure 5: Selection-corrected price indices per style

This figure shows selection-corrected price indices for each style classification, normalized to an index value of 100 in 1972. *Top 100* refers to the index of paintings by top 100 artists based on the total value of sales (in U.S. dollars) of all paintings by the artist over the sample period.

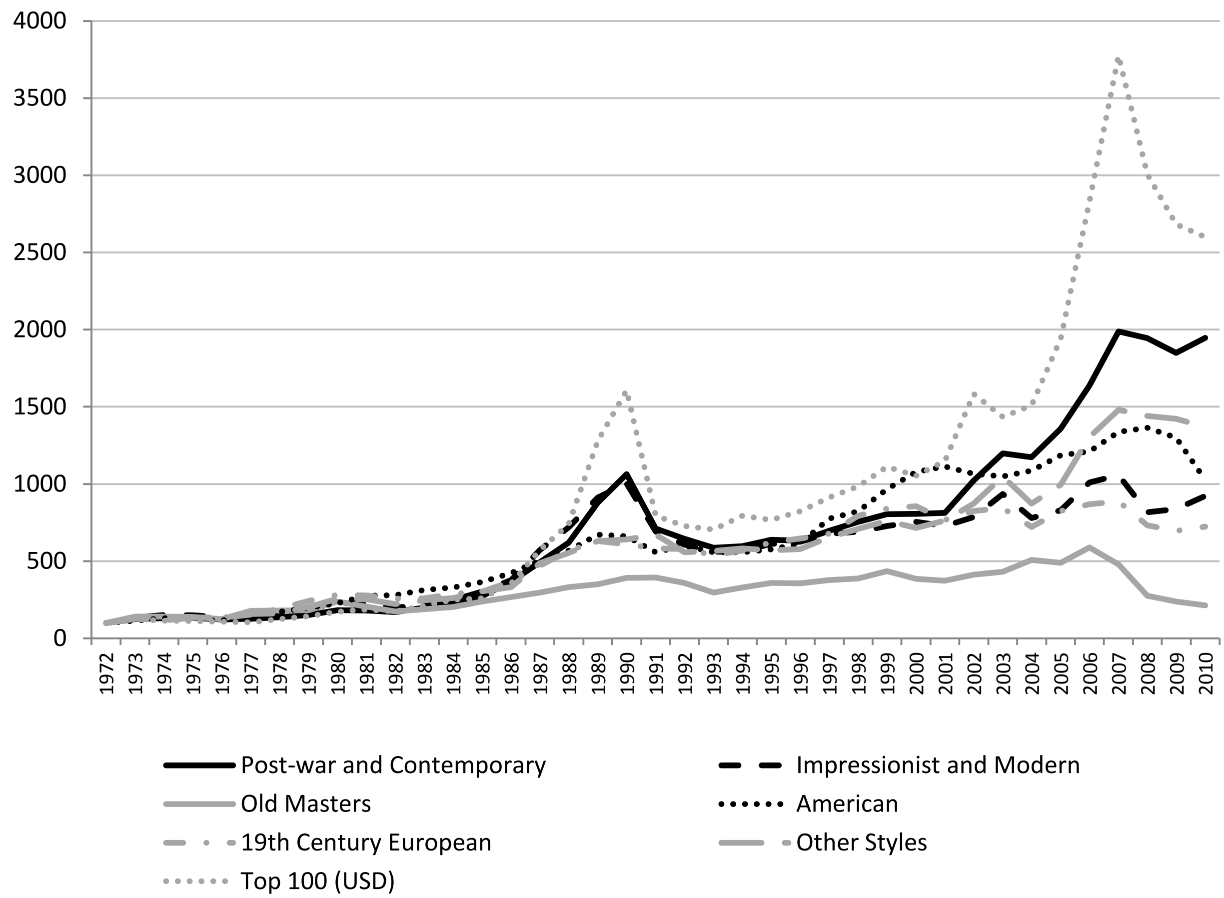

Figure 6: Annualized sale-to-sale returns by return horizon

This figure shows the relation between the logarithm of annualized sale-to-sale returns (on the left-hand vertical axis) and the time between sales (in years, on the horizontal axis) in the repeat-sales data. The solid and striped lines represent the average and the median annualized log return between observed sales, respectively. The vertical bars are the number of observations in each bin, measured on the right-hand vertical axis.

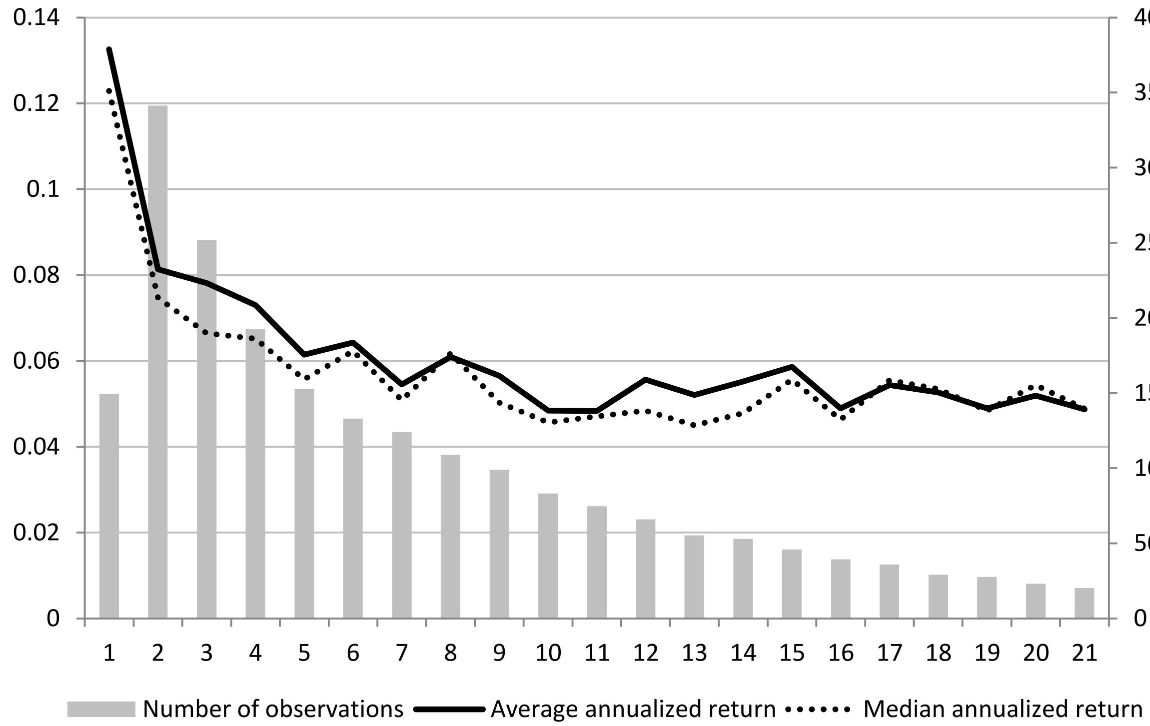

Figure 7: Time series of annualized returns and trading intensity

This figure graphs the time series of the average annualized log sale-to-sale returns (represented on the left-hand vertical axis) and trading intensity (on the right-hand axis). The average log sale-to-sale return is computed over consecutive sales of paintings for which the second sale falls in the given year. Trading intensity is calculated as the number of sales in a calendar year as a percentage of all sales over the 1980 to 2010 period.

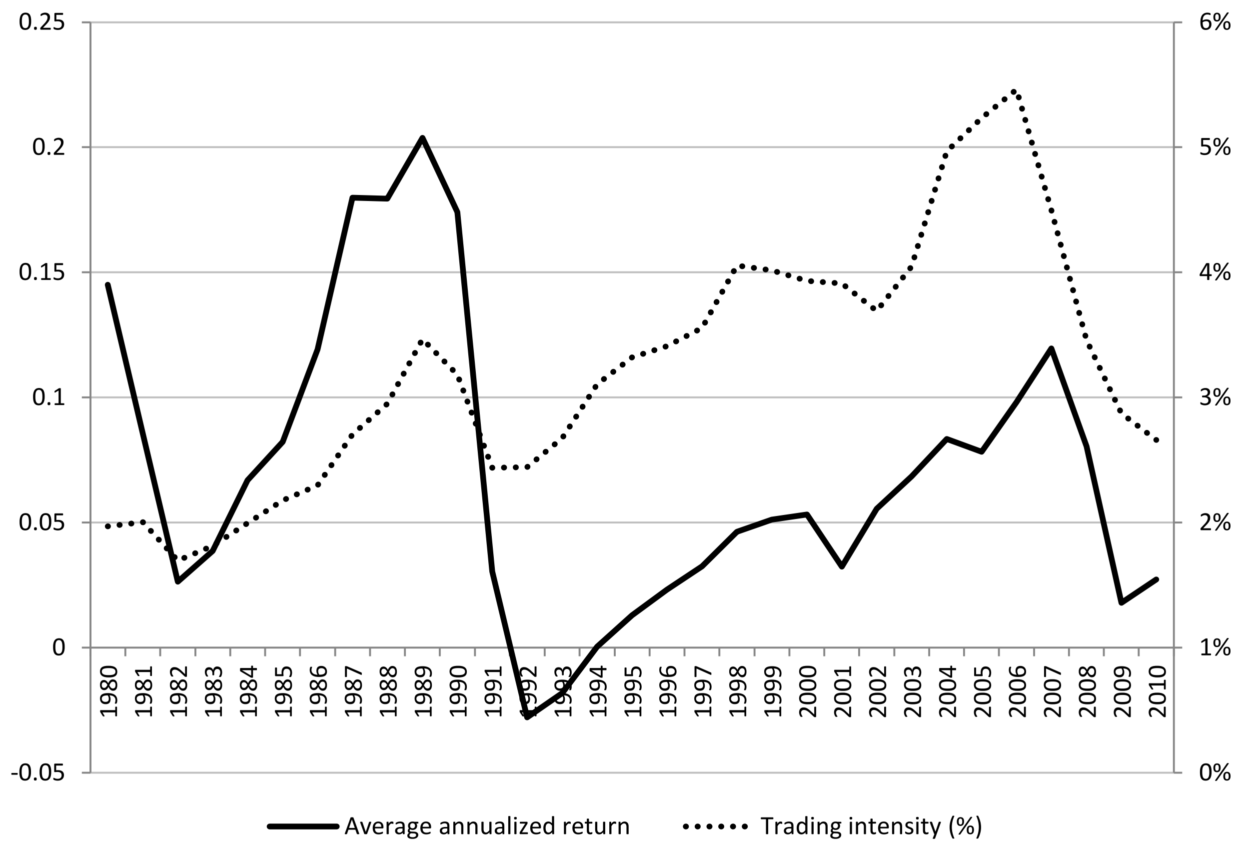

Table 1: Summary statistics

This table reports summary statistics for the sample of paintings in the Blouin Art Sales Index (BASI) dataset from 1972 to 2010. Panel A presents descriptive statistics for the repeat-sales sample that contains paintings that sold at least twice during the sample period (left columns), and the full BASI dataset (right columns). The unit of observation is a sale of a painting at auction. *Hammer price* is the auction price in thousands of U.S. dollars. *Surface* is the surface of the painting in thousands of squared millimetres. *Deceased < 2 yrs* is a dummy variable equal to one when the sale occurs within two years after the artist deceases, and zero otherwise. *Christie's* and *Sotheby's* are dummy variables that equal one if the painting is auctioned at Christie's or Sotheby's, respectively, and *London* and *New York* are dummy variables that equal one if the painting is auctioned in London or New York, respectively. *Top 100 Artists* is a dummy variable equal to one when the artist is in the top 100 in terms of total value of sales (in U.S. dollars) over 1972 to 2010, and zero otherwise. The remaining variables in Panel A represent style classifications. The last column shows test statistics for the difference in means *t*-statistics (for Hammer price and Surface) and the difference in proportions *z*-statistic (for the other variables) between the full and the repeat-sales samples. The sale-to-sale returns in Panel B are for the repeat-sales sample only, and calculated as the natural logarithm of the ratio of the current and prior hammer price of a painting. ***, ** and * indicate statistical significance at the 1, 5 and 10 percent level, respectively.

Panel A. Descriptive statistics for the full and repeat-sales sample

	Repeat-sales sample (42,548 sales)			Full sample (2,302,738 sales)			Difference statistic
	Mean	Median	St. Dev.	Mean	Median	St. Dev.	
Hammer price ($000s)	61.9	6.2	444.2	28.6	3.0	395.5	-15.39***
Surface	546.9	331.0	792.1	491.8	306.5	811.8	-14.19***
Deceased < 2yrs	2.12%			1.59%			-8.71***
Christie's	21.54%			15.57%			-33.59***
Sotheby's	25.46%			15.82%			-53.75***
London	19.47%			14.15%			-31.07***
New York	21.12%			9.75%			-77.66***
Top 100 Artists	9.62%			2.77%			-83.52***
Post-war and Contemporary	15.81%			11.75%			-25.69***
Impressionist and Modern	34.03%			22.32%			-57.29***
Old Masters	5.32%			9.87%			31.26***
American	11.57%			7.29%			-33.47***
European 19th Century	23.83%			31.28%			32.89***
Other Style	9.44%			17.49%			43.46***

Panel B. Sale-to-sale returns for the repeat-sales sample (22,010 returns for 20,538 paintings)

	Mean	Median	St. Dev.
Sale-to sale return			
Arithmetic return	123.50%	42.37%	368.54%
Log return	43.94%	35.33%	78.08%
Years between sales	7.61	5.55	6.28
Annualized sale-to-sale return			
Arithmetic return	16.50%	7.53%	32.68%
Log return	6.90%	5.65%	16.69%

Table 2: Selection equation coefficients

Parameter estimates of three specifications of the selection equation (equation (3) in the text). *Log return* is the natural logarithm of the return since the prior sale of a painting. *Time* is the time in years since the prior sale. *Log surface* is the natural logarithm of the painting's surface in thousands of mm^2. *World GDP growth* is the yearly increase in worldwide GDP, obtained from Historical Statistics of the World Economy. The other variables are as defined in Table 1. *Sigma* is the standard deviation of the error term in equation (1). Standard errors are in parentheses. ***, ** and * indicate statistical significance at the 1, 5 and 10 percent level, respectively.

	A	B	C
Log return	0.375***	0.372***	
	(0.008)	(0.008)	
1972-1989			0.300***
			(0.014)
1990-1999			0.278***
			(0.009)
2000-2006			0.581***
			(0.013)
2007-2010			0.772***
			(0.027)
Time (yrs)	-0.401***	-0.033***	-0.022***
	(0.019)	(0.002)	(0.002)
Time squared	0.090***	0.001***	0.000***
	(0.009)	(0.000)	(0.000)
Log surface		0.005	0.002
		(0.007)	(0.008)
Deceased < 2yrs		0.104***	0.094***
		(0.031)	(0.031)
World GDP growth		-1.008***	-1.273***
		(0.253)	(0.309)
Intercept	-1.517***	-1.512***	
	(0.007)	(0.042)	
1972-1989			-1.487***
			(0.044)
1990-1999			-1.490***
			(0.044)
2000-2006			-1.518***
			(0.045)
2007-2010			-1.660***
			(0.046)
Sigma	0.282***	0.282***	0.288***
	(0.001)	(0.001)	(0.001)

Table 3: Relation between annualized returns and market shares by style

Panel A shows the correlation between the yearly market share of each style and the mean (left column) and median (right column) annualized sale-to-sale return for the style, computed over the returns for which the second sale falls in the given year. The market share of a specific style is the total sales of this style in a given year relative to all sales in that year, calculated from the full BASI dataset. Panel B shows the coefficients of a regression of the yearly style market shares on the annualized sale-to-sale returns by style, and dummy variables representing the styles (*Other Style* is the omitted variable). Standard errors are in parentheses. ***, ** and * indicate statistical significance at the 1, 5 and 10 percent level, respectively.

Panel A. Correlation coefficients between market shares and returns per style

	Mean annualized sale-to-sale return	Median annualized sale-to-sale return
Post-war and Contemporary	0.235	0.257
Impressionist and Modern	0.411	0.426
Old Masters	0.083	-0.053
American	0.391	0.467
European 19th Century	0.224	0.295
Other Style	-0.080	-0.178
Top 100 Artists	0.373	0.296

Panel B. Regression analysis (Dependent variable = Yearly market share by style)

	I	II
Mean annualized sale-to-sale return	0.116***	
	(0.035)	
Median annualized sale-to-sale return		0.130***
		(0.034)
Post-war and Contemporary dummy	0.069***	0.069***
	(0.009)	(0.009)
Impressionist and Modern dummy	0.240***	0.240***
	(0.009)	(0.009)
Old Masters dummy	-0.033***	-0.033***
	(0.009)	(0.009)
American dummy	0.019**	0.019**
	(0.009)	(0.009)
European 19th Century dummy	0.145***	0.144***
	(0.009)	(0.009)
Top 100 Artists dummy	-0.001	-0.001
	(0.009)	(0.009)
Intercept	0.086***	0.087***
	(0.007)	(0.007)
Adjusted R^2	83.7%	83.7%
Number of observations	264	264

Table 4: Summary statistics of annual index returns, 1972-2010

Panel A reports descriptive statistics of the annual returns to indices of paintings and stocks, and to one-month U.S. Treasuries over the period 1972 to 2010. *GLS* is the standard repeat-sales arithmetic index of paintings as estimated by Generalized Least Squares. The *MCMC* index is the non-selection-corrected index from our Markov chain Monte Carlo estimator. The selection-corrected art indices *A* through *C* are as described in Table 2. The stock index is the CRSP value-weighted stock index return, including distributions. The left columns show results for arithmetic returns, and the right columns show log returns. S.R. stands for the annual Sharpe Ratio. Panel B reports the correlation coefficients between excess returns on the stock index and the various art indices.

Panel A. Descriptive statistics of annual index returns

	Arithmetic returns			Log returns		
	Mean	St. dev.	S.R.	Mean	St. dev.	S.R.
Returns on non-selection corrected art indices:						
GLS	11.13%	12.84%	0.391	9.88%	11.77%	0.340
MCMC	11.10%	12.71%	0.393	9.88%	11.66%	0.342
Returns on selection-corrected art indices:						
A	7.29%	12.72%	0.096	6.34%	11.96%	0.041
B	7.33%	12.71%	0.099	6.38%	11.96%	0.045
C	6.70%	13.29%	0.047	5.72%	12.52%	-0.011
Returns on other assets:						
Stocks	11.70%	19.01%	0.297	9.45%	18.83%	0.192
Treasuries	6.09%	3.34%	-	5.86%	3.12%	-

Panel B. Correlation coefficients between annual excess returns on stocks and art indices

	Arithmetic returns	Log returns
GLS	-0.075	-0.053
MCMC	-0.067	-0.045
A	-0.022	0.008
B	-0.019	0.012
C	-0.011	0.015

Table 5: Optimal asset allocation

Panel A shows the mean-variance tangency portfolio weights on paintings (using the non-selection corrected GLS index, and the selection-corrected indices A, B, and C from Table 2) and stocks (the CRSP value-weighted index including distributions), based on arithmetic (left columns) or log returns (right columns). Panels B, C and D show the optimal weights for a one-period mean-variance utility investor with risk aversion (γ) equal to two, five, and ten, respectively. Short sales are not allowed. Returns are measured over 1972 to 2010. Sharpe Ratios are annual.

	Arithmetic returns				Log returns			
	GLS	selection-corrected			GLS	selection-corrected		
		A	B	C		A	B	C
Panel A. Tangency portfolio weights								
Paintings	0.649	0.342	0.346	0.198	0.724	0.248	0.260	0.000
Stocks	0.351	0.658	0.654	0.802	0.276	0.752	0.740	1.000
Sharpe Ratio	0.510	0.314	0.315	0.301	0.400	0.196	0.197	0.192
Panel B. Mean-variance utility, risk aversion γ = 2								
Paintings	0.625	0.269	0.273	0.169	0.753	0.143	0.154	0.000
Stocks	0.375	0.731	0.727	0.768	0.247	0.503	0.503	0.506
Treasuries	0.000	0.000	0.000	0.063	0.000	0.354	0.343	0.494
Sharpe Ratio	0.509	0.313	0.314	0.301	0.399	0.196	0.197	0.192
Panel C. Mean-variance utility, risk aversion γ = 5								
Paintings	0.648	0.148	0.152	0.061	0.595	0.057	0.062	0.000
Stocks	0.340	0.310	0.310	0.310	0.218	0.201	0.201	0.203
Treasuries	0.012	0.542	0.538	0.629	0.187	0.742	0.737	0.797
Sharpe Ratio	0.510	0.314	0.315	0.301	0.400	0.196	0.197	0.192
Panel D. Mean-variance utility, risk aversion γ = 10								
Paintings	0.322	0.074	0.076	0.031	0.297	0.029	0.031	0.000
Stocks	0.170	0.154	0.154	0.155	0.108	0.100	0.100	0.100
Treasuries	0.508	0.772	0.770	0.815	0.595	0.871	0.869	0.900
Sharpe Ratio	0.510	0.314	0.315	0.301	0.400	0.196	0.197	0.192

Table 6: Optimal asset allocation with different painting styles

Panel A shows mean-variance tangency portfolio weights (in the column *Tang ptf*), and the optimal weights for a one-period mean-variance utility investor with risk aversion (γ) equal to two, five, and ten, on paintings of different styles and public stocks (the CRSP value-weighted index including distributions). The left columns use arithmetic returns, the right columns on log returns. Panel B also includes an index for Top 100 artists (based on total value of sales over the sample period). Short sales are not allowed. Returns are measured over 1972 to 2010. Sharpe Ratios are annual.

Panel A. Style indices

	Arithmetic returns				Log returns			
	Tang	Risk aversion (γ)			Tang	Risk aversion (γ)		
	ptf	2	5	10	ptf	2	5	10
Post-war and Contemporary	0.391	0.410	0.308	0.154	0.504	0.495	0.217	0.108
Impressionist and Modern	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Old Masters	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
American	0.146	0.000	0.007	0.003	0.000	0.000	0.000	0.000
European 19th Century	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Other Styles	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Stocks	0.464	0.590	0.327	0.163	0.496	0.505	0.213	0.107
Treasuries	-	0.000	0.359	0.680	-	0.000	0.570	0.785
Sharpe Ratio	0.380	0.373	0.377	0.377	0.247	0.247	0.247	0.247

Panel B. Styles indices and index for Top 100 Artists

	Arithmetic returns				Log returns			
	Tang	Risk aversion (γ)			Tang	Risk aversion (γ)		
	ptf	2	5	10	ptf	2	5	10
Post-war and Contemporary	0.057	0.000	0.050	0.024	0.497	0.495	0.217	0.108
Impressionist and Modern	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Old Masters	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
American	0.107	0.000	0.000	0.000	0.000	0.000	0.000	0.000
European 19th Century	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Other Styles	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Top 100 Artists	0.303	0.395	0.189	0.096	0.006	0.000	0.000	0.000
Stocks	0.533	0.605	0.327	0.164	0.497	0.505	0.213	0.107
Treasuries	-	0.000	0.434	0.715	-	0.000	0.570	0.785
Sharpe Ratio	0.394	0.393	0.393	0.393	0.247	0.247	0.247	0.247

Table 7: Optimal asset allocation with Dimson correction

Panel A shows mean-variance tangency portfolio weights on paintings (using the non-selection corrected GLS index, and the selection-corrected indices A, B, and C from Table 2) and stocks (the CRSP value-weighted index including distributions), based on arithmetic (left columns) or log returns (right columns). Panels B, C and D show optimal weights for a one-period mean-variance utility investor with risk aversion (γ) of two, five, and ten, respectively. The covariance matrix of (excess) returns is computed using the Dimson (1979) correction with leads and lags of one year. Short sales are not allowed. Returns are measured over 1972 to 2010. Sharpe Ratios are annual.

	Arithmetic returns				Log returns			
	GLS	selection-corrected			GLS	selection-corrected		
		A	B	C		A	B	C
Panel A. Tangency portfolio weights								
Paintings	0.476	0.264	0.263	0.147	0.573	0.207	0.207	0.000
Stocks	0.524	0.736	0.737	0.853	0.427	0.793	0.793	1.000
Sharpe Ratio	0.455	0.340	0.340	0.326	0.333	0.208	0.208	0.203
Panel B. Mean-variance utility, risk aversion γ = 2								
Paintings	0.463	0.171	0.172	0.055	0.597	0.085	0.086	0.000
Stocks	0.537	0.829	0.828	0.876	0.403	0.540	0.539	0.538
Treasuries	0.000	0.000	0.000	0.069	0.000	0.376	0.375	0.462
Sharpe Ratio	0.455	0.337	0.337	0.325	0.333	0.207	0.207	0.203
Panel C. Mean-variance utility, risk aversion γ = 5								
Paintings	0.352	0.094	0.095	0.022	0.316	0.034	0.035	0.000
Stocks	0.372	0.355	0.354	0.351	0.221	0.216	0.216	0.216
Treasuries	0.276	0.551	0.551	0.626	0.463	0.750	0.750	0.784
Sharpe Ratio	0.455	0.339	0.339	0.325	0.333	0.207	0.207	0.203
Panel D. Mean-variance utility, risk aversion γ = 10								
Paintings	0.174	0.047	0.047	0.011	0.158	0.017	0.017	0.000
Stocks	0.187	0.177	0.177	0.175	0.110	0.108	0.108	0.108
Treasuries	0.638	0.775	0.776	0.814	0.732	0.875	0.875	0.892
Sharpe Ratio	0.455	0.339	0.339	0.325	0.333	0.207	0.207	0.203