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 that trade endogenously. Using a sample of 32,928 paintings that sold repeatedly between 1960 and 2013, we find an asymmetric V-shaped relation between sale probabilities and returns. Adjusting for the resulting selection bias cuts average annual index returns from 8.7 percent to 6.3 percent, lowers Sharpe ratios from 0.27 to 0.11, and materially impacts portfolio allocations. Investing in a broad portfolio of paintings is not attractive, but targeting specific styles or top-selling artists may add value. The methodology extends naturally to other asset classes.

JEL classification: D44, G11, Z11

Keywords: Art investing; Selection bias; Portfolio allocation; MCMC

Current Draft: June 25, 2015

* Arthur Korteweg (korteweg@marshall.usc.edu) is from the University of Southern California Marshall School of Business, Roman Kräussl (roman.kraussl@uni.lu) is from the 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, and University of Glasgow. We thank Andrew Karolyi (the editor), two anonymous referees, and Andrew Ang, Stephen Brown, Les Coleman, Elroy Dimson, Will Goetzmann, Bruce Grundy, Chad Kendall, Bryan Lim, Stefano Lovo, Andrew Metrick, George Parker, Caroline Sasseville, Christophe Spaenjers, Mark Westerfield, and participants at the 2013 European Finance Association meetings, and the 2014 "Art: Minds and Markets" conference at Yale University for helpful comments. We are especially grateful to Elroy Dimson, Paul Marsh, and Mike Staunton for sharing their global equities and bonds data.

Over the last three decades investors have started allocating increasingly larger shares of their portfolios to alternative assets. Many of these alternative asset classes, such as private equity and real estate, and even certain traditional assets such as corporate bonds, are highly illiquid, complicating return evaluation. In particular, when sales are endogenously related to the performance of the asset, a sample selection problem arises that is pervasive across asset classes. This paper develops a methodology to quantify the magnitude of the selection bias, and demonstrates its empirical first-order importance when evaluating investment performance and constructing optimal portfolios that include alternative assets.

Since the turn of the millennium, paintings (and other collectibles) have garnered increasing interest among alternative asset investors. Fueled by a strong rise in global wealth (and high net worth individuals in particular), and the search for yield in an environment of low interest rates and stock returns, the art[1] auction market doubled in sales volume between 2002 and 2013. By 2013, the global art market, though considerably smaller than real estate, was at least as large as the venture capital market in terms of sales and assets under management.[2] In recent years, institutional funds have sprung up to allow investors access to diversified investments in art, and by 2013, there were 104 such art funds in existence.[3] According to the Art & Finance Report 2014, a joint publication by Deloitte Luxembourg and ArtTactic, wealthy investors allocate 6 to 18 percent of their total wealth to art and collectibles (depending on the

[1] Following the literature, we use the terms "art" and "paintings" interchangeably throughout the paper.

[2] Global art sales have grown strongly in the previous decade, and by 2013 amounted to EUR 47 billion (McAndrew, 2014). In comparison, the National Association of Realtors reported 2013 existing home sales of $1.2 trillion, and the U.S. Census reported new home sales of $139 billion in the U.S. alone. The global venture capital industry invested around $47 billion per year between 2006 and 2013, with little to no growth (Pearce, 2014). In 2013, the U.S. VC industry had $193 billion under management in (NVCA Yearbook 2015). Though it is difficult to get a good estimate of the total value of paintings, the art stored in Geneva's Freeport alone is estimated to be worth around $100 billion (The Economist, 2013).

[3] Examples of art fund managers are the Fine Art Fund Group, Anthea Art Investments, the Art Vantage Fund, and the Artemundi Global Fund. The first has $200 million under management, the others $15 to 40 million each.

region), and the majority of wealth managers and family offices strongly believe that there is a role for art in wealth management (Picinati di Torcello and Petterson, 2014).

The role of art in investments is still a hotly debated topic among practitioners and academics. While the academic literature has traditionally found that the returns on art are lower than those on stocks (see Frey and Eichenberger, 1995, and Burton and Jacobsen, 1999 for an overview), recent studies have found that there may be value in including art in investment portfolios, partly due to its low - or even negative - correlation with other asset classes (see the review in Ashenfelter and Graddy, 2003, and recent work by Mei and Moses, 2002, and Taylor and Coleman, 2011). Renneboog and Spaenjers (2013) find that although the Sharpe ratio for art does not surpass that for stocks and bonds, it is higher than for other popular alternative asset classes, such as commodities and real estate. Moreover, art has performed very well in recent years, as underscored by the greater than 100% increase in the popular Mei and Moses index between 2002 and 2013, a time period that is excluded from the earlier academic studies.

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 sells. Goetzmann (1993, 1996) argues that these sales are endogenous, and 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 sales, it is possible to observe high and positive returns even though overall values of paintings are declining. Consistent with this notion, the realization utility model of Barberis and Xiong (2012) predicts that large gains are indeed more likely to realize than small ones. However, in their model, losses are realized purely due to random liquidity shocks, and the probability of a sale is unrelated to the size of the loss. Another set of theories predicts that the probability of sale increases in the size of the loss, resulting in a V-shaped relation between sale probabilities and returns. For example, this happens if the incidence of liquidity shocks is correlated with the size of the loss[4], or if people trade off incurring a loss with a reset of their reference points (Ingersoll and Jin, 2013). On the other hand, Meng (2014) predicts that loss aversion can lead to the exact opposite, i.e., an inverse V-shaped relation. Quantifying the relation between returns and sale probabilities in illiquid assets, and the direction and magnitude of the resulting selection bias are therefore important empirical questions that have thus far proven difficult to answer, in large part due to lack of empirical methods.

We present a new and flexible econometric model of art indices, based on the framework developed by Korteweg and Sorensen (2010, 2014), that generalizes the standard repeat sales regression (RSR; see Bailey, Muth, and Nourse, 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 potentially unobserved valuations and returns between sales, as well as the probability of observing a sale 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 sold. We also use the price information in failed auctions (called "buy-ins"), which are typically ignored in the literature.

Using Bayesian MCMC methods, we estimate the model on a new proprietary auction data set from which we construct the largest sample of repeat sales of paintings in the literature to date, with 32,928 paintings being sold a total of 69,103 times between 1960 and 2013. To our knowledge, ours is the first paper to use this data set.

[4] A similar phenomenon may have happened in residential real estate during the collapse of the housing bubble during the late 2006 to 2008 period, when a disproportional fraction of the housing transactions were foreclosures (e.g., Campbell, Giglio and Pathak, 2011), whose drop in value were likely not representative for the aggregate housing market.

We find that the relation between the change in the value of paintings since purchase and the probability of a sale is V-shaped, such that large gains are more likely to realize than small gains, and large losses are more likely than small ones. Ben-David and Hirshleifer (2012) also find a V-shaped relation for stock sales by retail investors, but unlike their results, we find a significant jump in the probability of a sale when a painting turns from a marginal loss to a marginal gain. This implies that gains are usually more likely to realize than losses, consistent with the disposition effect as hypothesized by Goetzmann (1993, 1996).

Our results show that selection bias is of first-order economic importance. 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 1960, the RSR index is 5,429 in 2013, the end of our sample period, whereas the selection-corrected index ends around 1,895. This implies that the average annual return on the corrected index is 6.3 percent, which is 28 percent lower than the 8.7 percent average return on the uncorrected RSR index. The annual Sharpe ratio drops nearly 60 percent, from 0.27 to 0.11.

The strength of the selection bias varies strongly in both the time-series and in the cross-section. The bias may even reverse in times of large art market downturns, when large losses are overrepresented in the data due to the V-shaped relation between returns and sale events. This also causes excess volatility in the uncorrected index. We find evidence of both these effects in the data. In the cross-section, the more popular painting styles typically sell for smaller gains while people wait for higher gains before selling unpopular styles. This is consistent with increased speculative selling of popular paintings (conform the model by Lovo and Spaenjers, 2014), and amplifies the selection bias on gains.

Selection bias has important implications for asset allocation decisions. Over our sample period, a mean-variance investor who ignores selection bias would allocate a portfolio weight of one quarter to one third to a broadly diversified portfolio of art that aims to track the aggregate art market (with the remainder allocated to global stocks, corporate bonds, real estate, and commodities). The investor thinks that this portfolio has a Sharpe ratio of 0.64. However, since art returns are subject to selection bias, the Sharpe ratio is in fact only 0.50, which is 22 percent less than the investor's perception, and 14 percent lower than the 0.59 Sharpe ratio of a portfolio that excludes art altogether. Correcting for selection bias, the investor should optimally forego investing in a broad art portfolio, barring substantial non-monetary utility from owning and enjoying art (Mandel, 2009).[5] This result is robust to the effects of illiquidity, transaction costs, and higher moments of art returns on portfolio allocations. We do find suggestive evidence that there may be value in following a targeted strategy aimed at a particular style, or at top-selling artists.

This is not the first paper to consider sample selection in illiquid assets. The issue was first raised in the real estate literature by Case, Pollakowski and Wachter (1991) and Haurin and Hendershott (1991). Subsequent work used standard Heckman models to estimate the bias (e.g., Jud and Seaks, 1994, Gatzlaff and Haurin, 1997, 1998, Munneke and Slade, 2000, 2001, Hwang and Quigley, 2004, and Goetzmann and Peng, 2006), an approach that was also adopted for other assets such as venture capital (Hwang, Quigley and Woodward, 2006) and art (Collins, Scorcu, and Zanola, 2007, and Zanola, 2007). However, the Heckman approach is problematic as it ignores the underlying dynamics of the price and selection processes. In contrast, our model

[5] If investing in the aggregate art market through a fund (a more feasible strategy for investors who cannot afford a diversified exposure to the art market through the purchase of individual, often expensive paintings), the investor experiences no consumption utility of owning art due to lack of access to the artworks held by the fund.

produces the unconditional price path of all paintings, which is useful to quantify the size of the selection bias in returns, among other applications.

Our paper is most closely related to Cochrane (2005) and Korteweg and Sorensen (2010), who consider venture capital returns, and Korteweg and Sorensen (2014), who examine the distribution of loan-to-value ratios in residential real estate. These papers also impose a selection equation on top of the price process. However, the approach in the Cochrane and Korteweg-Sorensen papers imposes a linear and continuous relation between returns and the selection variable that determines sale probabilities. Our model is more flexible, allowing for both non-linearities and discontinuities. Our results show that these features are economically important, providing a more realistic and more nuanced picture of selection bias, including its time-series and cross-sectional variation. In addition, we exploit the information in failed auctions, and we show how to estimate separate style indices. The above papers also do not consider optimal portfolio allocations.

The methodology developed in this paper extends naturally to other illiquid assets. For example, the model can easily accommodate structural corporate bond pricing equations, which are highly non-linear in the state variable, or an “under-water” effect in homeownership, where homeowners are unable to move if their home value drops below the outstanding mortgage balance, causing a discontinuity in sales probabilities. The information in unsold real estate listings is akin to buy-ins, and can be taken into account in the model. Uncovering these relations in other asset classes and quantifying their impact on the size and direction of the bias in returns is an important avenue for future work.

1. Art market and price data

Following the extant art literature we use auction prices to construct price indices. Many paintings sell in public auctions, and auction data is more reliable and more easily obtained than private-market data (i.e., gallery or direct-from-artist prices). It is generally accepted that auction prices set a benchmark that is used in the private market (Renneboog and Spaenjers, 2013).

Our auction data come from the Blouin Art Sales Index (BASI), an independent database of artworks sold at over 350 auction houses worldwide, including Christie's and Sotheby's. BASI sources its data from Hislop's Art Sales Index, the primary source of price information in the world of fine art, supplemented with catalogue data from auction houses (both electronic and hard copy). BASI is presently the largest known database of artworks, containing roughly 4.6 million works of art (of which 2.3 million are paintings) by more than 225,000 individual artists over the period 1922 to 2013.

For each sold painting in our data set, we have detailed information about the painting, the artist, and the auction. We know the painting's title, artist, year of creation, size, whether it was signed or stamped by the artist, and its medium (e.g., "oil on canvas", or "acrylic on board"). The database also categorizes each painting into one of six main styles as defined by Christie's and Sotheby's: Post-war and Contemporary, Impressionist and Modern, Old Masters, American, 19th Century European, and a residual "Other" style category. For each artist, we observe their name, nationality, year of birth, and year of death (where applicable). We also know the date of the auction, and the auction house at which the painting was brought up for sale.

Prior to the auction, auction house experts provide high and low price estimates between which they expect the painting to sell. Sellers set a reserve price (usually called "reservation price" in economics), which is the lowest price the seller is willing to accept. Auction houses do not disclose the reserve price, but auction rules require it to be set at or below the low price estimate. The auction is conducted as an ascending bid (i.e., English-style) auction, where the auctioneer calls out increasingly higher prices. When a bid is solicited that no other bidder is prepared to exceed, the auctioneer strikes his hammer, and - provided it exceeds the reserve price - the painting is sold at this highest bid price (called the "hammer price"). In our data, we convert all hammer prices and price estimates to U.S. dollars using the spot rate at the time of sale. The hammer price excludes any transaction costs that the buyer pays to the auction house.

If the highest bid did not surpass the reserve price, the painting remains unsold, and it is said to be "bought in".[6] Starting in 2007 we observe when a painting is bought in, and we have information about the painting, the artist, and the auction event. Beggs and Graddy (2008) study 43 paintings that were bought in, and find that they experience lower future returns than paintings without a buy-in. Since buy-ins contain price information, we will use them in some of our estimates.

Our price index methodology relies on paintings that sell repeatedly. We construct a sample of repeat sales by matching auction sales records using artists' names, artwork names, painting size, medium, and signature (similar to the matching procedures in Taylor and Coleman, 2011, and Renneboog and Spaenjers, 2013). The sample period, in which we observe a sufficient number of repeat sales to construct a price index, ranges from 1960 to 2013.

To eliminate false matches, we remove paintings from the same artist with the generic titles "untitled" and "landscape", which represent 1.9% of paintings. We extensively checked whether the remaining potential repeat sales are true repeat sales by manually searching for the painting's "provenance". The provenance is the documented history of an artwork's creation and

[6] Auction firms generally charge sellers a percentage of their reserve price in case the artwork goes unsold. This encourages sellers to keep their reserve price low (Ashenfelter and Graddy, 2003).

ownership, and may be found in auction catalogues, the online databases Artnet and Artprice, and on the websites of auction houses.[7] When we are in doubt about whether we are dealing with a true repeat sale, we delete it from our sample. Moreover, dealers sometimes change titles of works to a generic title to curtail the ability of buyers to find prior auction prices, which may cause us to miss some true repeat sales. Our final sample includes 69,103 sales of 32,928 unique paintings. Appendix A contains a comparison between the repeat sales sample and the full BASI data set.

Figure 1 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.[8]

[Please insert Figure 1 here]

Table 1 Panel A shows descriptive statistics of the paintings. The average (median) hammer price is $150,600 ($14,500), with a long right tail of extremely expensive paintings. The average (median) low price estimate is 86 percent (82 percent) of the hammer price, and the average (median) high estimate is 119 percent (114 percent). The size of an average painting is 630,300mm^2, or 0.63m^2. Nearly two percent of sales occur within two years after the artist has deceased. Around 33 percent of sales take place at Christie's and Sotheby's auction houses each. For 27 percent of sales, the auction house is located in London, and 29 percent are sold in New York. Of the major styles, Impressionist paintings are sold most often in the sample, accounting for 26 percent of all sales, and American paintings are sold the least, making up 10 percent of sales. About 15 percent of sold paintings are by artists with total dollar sales in the top 100 of the full BASI database over the preceding decade.

[Please insert Table 1 here]

---

[7] For example, Christie's and Sotheby's provide online provenance information on all auction sales since 1998.
[8] We observe 229 paintings with four sales, 18 paintings that sold five times, and one painting sold six times.

The differences in sales frequencies across styles is in part due to styles' relative abundance in the population, but there is also time variation in the relative sale frequencies due to time-variation in the relative popularity of styles. We return to this in Section 3.

Between 2007 and 2013 the paintings in our repeat sales sample went up for auction an additional 3,854 times (representing 20.6% of the total number of auctions over the period), but these auctions resulted in a buy-in. Appendix A contains more detailed descriptive statistics of the buy-ins.

Table 1 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 109.4 percent on average. The median return is 47.1 percent, and the standard deviation is 187.4 percent. With an average time between sales of 9.4 years, this translates to an average (median) annualized return of 12.4 percent (6.7 percent), with a standard deviation of 21.9 percent. Negative returns occur regularly, though annualized returns below –30 percent or above 70 percent are rare. Panel B also shows logarithmic returns, which are lower than the arithmetic returns. The average log return is 45.9 percent (5.6 percent annualized), with a median of 38.6 percent (4.8 percent annualized) and a standard deviation of 72.0 percent (13.5 percent annualized).

2. A selection model of prices and sales of artworks

2.1. Standard repeat sales regression model

We first introduce the classic repeat sales regression (RSR) model, which was originally developed by Bailey et al. (1963) and Case and Shiller (1987) to estimate real estate price indices, but has subsequently been modified and used for other illiquid asset classes, such as municipal bonds (Wilkoff, 2013), venture capital (e.g., Peng, 2001, Woodward and Hall, 2003, Korteweg and Sorensen, 2010), as well as art (e.g., Baumol, 1986, Buelens and Ginsburgh, 1993, Goetzmann, 1993, 1996, Mei and Moses, 2002, Zanola, 2007, and Goetzmann, Renneboog, and Spaenjers, 2011).

The standard RSR model decomposes 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*. The arithmetic price index is constructed from the time series of δ as shown in Goetzmann (1992) and Goetzmann and Peng (2002). Appendix B explains the derivation in more detail. The second component, $\varepsilon_i(t)$, is an idiosyncratic return that is particular to the individual artwork. Following the standard in the literature, 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, Table 1 shows that paintings sell infrequently, with an average time between sales of 9.4 years (median 7.3 years). 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:

$$r_i^h(t) = \sum_{\tau=t-h+1}^{t} \delta(\tau) + \varepsilon_i^h(t). \tag{2}$$

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 generalized least squares (GLS) regression techniques, scaling return observations by $1/\sqrt{h}$ to correct for heteroskedasticity.

2.2. Selection model

The δ estimates from the RSR model are consistent as long as the indicator variables in the 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 inflated 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, for example due to the disposition effect.[9]

To correct for selection bias we augment the RSR model with a selection equation that describes the probability of a sale,

$$w_i(t) = \sum_{k=1}^{K} g_k(p_i(t); Z_i(t)) \cdot \alpha_k + W_i(t)'\alpha_w + \eta_i(t). \tag{3}$$

A sale of artwork *i* at time *t* occurs whenever the latent variable $w_i(t)$ is greater than zero, and remains unsold otherwise. $Z_i(t)$ and $W_i(t)$ are sets of observed covariates, e.g., style fixed effects, characteristics of the artist, the painting, the prior sale price, or the state of the economy.

[9] The disposition effect states that investors are more likely to sell assets that have risen in price since purchase while holding on to those that have dropped in value, and has been documented in real estate (e.g., Genesove and Mayer, 2001), public equities (e.g., Shefrin and Statman, 1985, Odean, 1998, Grinblatt and Keloharju, 2001, Frazzini, 2006, and Ben-David and Hirshleifer, 2012), executive stock options (Heath, Huddart, and Lang, 1999), and online gambling (Hartzmark and Solomon, 2012).

The error term, $\eta_i(t)$, is i.i.d. normal with mean zero and variance normalized to one, and independent of $\varepsilon_i(t)$.[10]

The summation term in equation (3) captures selection effects that are possibly non-linearly related to price. The main specification in the empirical implementation is

$$\sum_{k=1}^{K} g_k(p_i(t); Z_i(t)) \cdot \alpha_k = I\left(r_i^p(t) \leq 0\right) \cdot \left[r_i^p(t)\alpha_1 + r_i^p(t)^2\alpha_2\right] \qquad (4)$$
$$+I\left(r_i^p(t) > 0\right) \cdot \left[\alpha_3 + r_i^p(t)\alpha_4 + r_i^p(t)^2\alpha_5\right],$$

where $r_i^p(t)$ is the (non-annualized) log return on the painting since its prior sale, and $I(x)$ is an indicator variable that equals one when the expression x is true, and zero otherwise.

This specification has three attractive features. First, α_3 allows for the possibility of a jump in the likelihood of a sale around $r^p = 0$. Pure sign preference (i.e., preferring a marginal gain over a marginal loss), which is a possible cause of the disposition effect, predicts that $\alpha_3 > 0$. Tax-loss selling, on the other hand, predicts $\alpha_3 < 0$, as marginal losses are more likely to be realized than marginal gains.[11]

Second, the linear-quadratic components in equation (4) link the probability of a sale to the magnitude of r^p. For example, realization utility theory generally predicts that for $r^p > 0$, the probability of a sale increases with the size of the gain (e.g., Barberis and Xiong, 2012). On the other hand, Meng (2014) predicts that loss aversion can yield the opposite result, i.e., small gains are more likely to be realized than large gains. Whether loss aversion is strong enough for this to occur continues to be debated in the behavioral finance literature (see, for example, Barberis, 2013), and is thus an important empirical question. The relation between w and r^p may be non-

---

[10] The normalization is necessary, but without loss of generality, because the parameters in equation (3) are only identified up to scale, as in a standard probit model.
[11] In the U.S., net capital gains from selling collectibles (including paintings) are taxed at a maximum rate of 28% rather than the lower rate on stocks held for more than a year.

linear, and the inclusion of squared r^p allows for more flexibility in the relation between r^p and the probability of a sale.

Third, selling behavior may be asymmetric in gains versus losses (i.e., α_1 and α_2 need not equal α_4 and α_5, respectively). Prospect theory implies an upward sloping sale probability curve for losses, i.e., small losses are more likely realized than large losses. This prediction is shared by Meng's (2014) model. In their realization utility model, Barberis and Xiong (2012) instead predict a flat relation between the size of losses and sales (i.e., $\alpha_1 = \alpha_2 = 0$), because losses are realized only due to randomly occurring liquidity shocks. Yet another set of theories predicts that large losses may be more likely to realize than small ones. One such theory is that liquidity shocks are positively correlated with the size of the loss (for example, if people are forced to sell their paintings because they suffered a large loss). Ingersoll and Jin (2013) generate the same prediction in a dynamic extension of the Barberis and Xiong model, where people trade off the disappointment of realizing a loss with a "reset" of the reference point when reinvesting, subject to transaction costs.[12] Another potential driver is tax-loss selling behavior in the presence of transaction costs. A final story is belief updating induced by the arrival of news that causes substantial gains or losses, which in turn drives sales activity (Ben-David and Hirshleifer, 2012).

The selection model, consisting of the observation equation (1) and the selection equation, (3), nests the classic RSR model: If all selection coefficients ($\alpha_1 \ldots \alpha_K$) equal zero, then sales occur for reasons unrelated to price, there is no selection bias, and we recover the standard RSR model. By estimating and testing the selection coefficients, we allow the data to speak to the direction and importance of selection bias.

[12] In public equities, where transaction costs are low, the Ingersoll and Jin (2013) model is subject to the criticism that investors can almost continuously reset their reference point by selling and repurchasing the same stock. In illiquid markets such as art, this theory has more bite as transaction costs are higher, and the same painting may not come up for sale again for years.

2.3. Estimation and interpretation

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 of a painting, but also about its past and future prices, even at times when the artwork does not sell. 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 artworks, 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. Moreover, since the log-prices enter the selection equation non-linearly, we cannot rely on standard Kalman filtering techniques to filter out the unobserved price paths. We use Markov chain Monte Carlo (MCMC), a Bayesian estimation technique, with a single-state updating Metropolis-Hastings step (Jacquier, Polson, and Rossi, 1994) to filter the latent prices.[13] Appendix C describes the estimation algorithm in detail.

In order to interpret the selection equation as a model of sales behavior rather than simply a reduced-form model to correct for selection bias, the variables in equation (3) should be in the art

[13] On a technical note, we specify the proposal density in the Metropolis-Hastings step in such a way that if the relation between r^p and w is truly linear, we recover the Kalman filter exactly, and the acceptance rate is thus 100%.

owner's information set at the time of the decision to put the painting up for auction. Although $r_i^p(t)$ is not known to the seller until a sale is realized, under the assumption that the seller has an unbiased signal of $r_i^p(t)$ - an assumption dating back to at least Grossman and Stiglitz (1980) - we can preserve the structural interpretation while using the filtered $r_i^p(t)$ to estimate the model. This works because under this assumption the measurement error (here, the noise in the seller's signal) is uncorrelated with the observed variable (r_i^p), and it is therefore not subject to errors-in-variables bias (Berkson, 1950).

2.4. Reserve prices and buy-ins

In the structural interpretation of the selection equation above, we implicitly assumed that if a seller decides to take the painting to auction, a sale is sure to occur. In the presence of reserve prices, however, this is not necessarily true. For example, if the noisy price signal received by art owners (which they use to decide whether to take the painting to auction) is too rosy, the seller may set a reserve price that is too high relative to the actual arm's-length transaction price that would have realized in the absence of a reserve, resulting in a buy-in.

We estimate a set of specifications that includes the information from buy-ins. If a buy-in occurs, we use the low price estimate from the auction as an upper bound on the value of the painting, since by auction rules reserve prices cannot be higher than this estimate. The details of the estimation algorithm are in Appendix C.

2.5. Relation to the literature

The problem of sample selection in price indices of illiquid assets was first raised in the real estate literature by Haurin and Hendershott (1991) and Case, Pollakowski and Wachter (1991).

Subsequent papers[14] turned to Heckman sample selection models, but are subject to the above-mentioned problem with the independence assumption of the Heckman model. The advantage of our approach compared to the Heckman corrections for sample selection in repeat sales models, is that our estimates are based on an explicit time-series model of the underlying valuation and selection processes, producing a selection adjustment that is consistent with this underlying model. The model produces the unconditional price path of all paintings, which can be used to assess risk and return, among other applications.

Our model is closest to the selection-correction approach for venture-capital backed firms in Cochrane (2005) and Korteweg and Sorensen (2010), and for real estate in Korteweg and Sorensen (2014). Our model generalizes the Korteweg and Sorensen (2010, 2014) model, which assumes that equation (3) is linear in log-prices (i.e., K=1 and $g_1(\cdot)$ is linear in $p_i(t)$). We find that the non-linearities that we introduce are economically and statistically important. Further extensions that we introduce are the use of the information in buy-ins, and the estimation of sub-indices of artworks, for example for different styles of paintings. To estimate sub-indices, we replace equation (1) with

$$r_i(t) = X_i' \cdot \delta(t) + \varepsilon_i(t), \tag{5}$$

where 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.

[14] These papers include Jud and Seaks (1994), Gatzlaff and Haurin (1997, 1998), Munneke and Slade (2000, 2001), Hwang and Quigley (2004), and Goetzmann and Peng (2006) for real estate, Hwang, Quigley and Woodward (2006) for venture capital, and Collins, Scorcu, and Zanola (2007) and Zanola (2007) for art.

3. Empirical results

In this section we first present initial evidence for sample selection, then we discuss our model estimates, and then we compare art indices with and without selection adjustments.

3.1. Initial evidence for selection

We first show suggestive evidence of the existence and strength of the selection problem in the raw data, without relying on the econometric machinery of the selection model. Without selection (i.e., when paintings sell for reasons unrelated to returns), there should be no systematic relation between returns and the probability of a sale. This is not true in the data: Table 2 Panel A shows that the correlation between the average (median) annualized return since prior sale for the paintings that sell in a given year and the proportion of sales that fall in the same year is 0.409 (0.252).[15] This is consistent with sample selection where a positive shock to the value of paintings results in more sales of paintings, with higher average realized returns, and vice versa for a negative shock.

[Please insert Table 2 here]

Panel B shows that this result also holds by painting style. For all but the "Other" painting style, the average (and median) annualized returns of a given style have correlations ranging from 0.182 to 0.503 with the market share of that style. Market share is defined as the number of paintings of the style that sold over the year relative to the total number of paintings that sold across all styles in the same year. "Other" paintings is the exception, with a correlation of 0.086 for average returns and -0.063 for median returns. The correlation is also low for top-selling artists, though still positive at 0.067 (0.043) for average (median) returns. Finally, Panel

[15] The correlations are similar if we use the probability of a sale instead of the proportion of sales (i.e., normalizing the paintings that sell in a given year by the total number of paintings instead of the total number of sales).

C shows that the relation between market share and sale-to-sale returns is highly statistically significant in a pooled regression analysis that controls for style fixed effects. We will return to these results below, when we consider the impact of relative market shares on selling behavior.

3.2. Selection model estimates

The selection model requires 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 linear-quadratic relation between the selection variable, w, and the log return since last sale, r^p, with a discontinuity at $r^p = 0$, as specified in equation (4). We also include the time since the last sale, both linearly and squared, in all specifications. The time since last sale helps to identify the model, as it affects the probability of a sale but not the idiosyncratic return ($\varepsilon_i(t)$ in equation (1)) of the artwork over the next period, based on the commonplace assumption that prices incorporate all available public information (including the date on which the painting was last sold). There are several possible channels through which the time since last sale may affect sale probabilities, predicting different signs of the relation. For example, people may not want to resell a painting a month after they bought it, but only consider a sale once a certain amount of time has passed. Auction houses are also very reluctant to sell the same piece in the next auction event. Alternatively, it is possible that the longer a person owns a painting, the less likely they are to sell it, for example because of a familiarity or endowment effect. For the purpose of identification, what matters is that there exists some relation between time since sale and probability of sale irrespective of the idiosyncratic return, not which channel dominates. The third common feature across all models is style fixed effects, capturing differences in baseline probabilities of sale for different styles.

[Please insert Table 3 here]

Table 3 Model A shows the estimated coefficients of the selection model that includes the log return and time variables. The first key result is that all coefficients on r_i^p are statistically different from zero, which shows that sample selection is statistically important.

[Please insert Figure 2 here]

To facilitate interpretation of the coefficients and their economic significance, Figure 2 plots the estimated relation between r_i^p and the probability of sale over the next year. The figure reveals the second key result: a V-shaped relation between returns and sale probabilities. Large gains are more likely to be realized than small gains, and large losses are more likely realized than small losses.

The third main result is the jump in the sale probability around zero returns, as α_3 in Model A is significantly larger than zero. In other words, the probability of a sale jumps up as marginal losses turn to marginal gains. Figure 2 shows that one year after a sale, a marginal gain sells with a probability of 5.14 percent, versus 3.99 percent for a marginal loss, a difference of 1.15 percent. At five and ten year horizons this difference is 1.06 and 0.97 percent, respectively.[16]

Most importantly, gains are typically more likely to be realized than losses. For example, Figure 2 shows that a painting that increased in value by 10 percent in the year since its prior sale, sells with a probability of 5.33 percent, whereas a painting that lost 10 percent has a sale probability of 4.54 percent. The higher sale probability of paintings that increase in value means that gains are more likely to appear in the data, and indices that do not correct for selection therefore overstate the price appreciation of the art market. One exception, however, is that in

[16] The jump in the sale probability around zero returns is somewhat surprising if investors have a noisy signal of the price at the time when they decide to bring the painting to auction. Non-mutually exclusive explanations are that sellers set the reserve price at the prior sale price, or that bidders (or the auctioneer) anchor on the prior sale price, so that it is more likely that a painting will sell at a small gain than a small loss.

times of large market downturns, large losses may occur that – due to the V-shaped sale probability function – are more likely to realize. As we discuss below, this causes time-variation in not only the strength but also the direction of the selection bias, and causes the non-corrected index to overstate the drop in aggregate prices when overall art prices experience large drops.[17]

To further illustrate economic significance, Figure 2 also shows the expected return and the top and bottom quartiles of the selection-corrected return distribution. The sale probability at the expected one-year return is 5.25 percent, while the top and bottom quartiles are 5.56 and 4.38 percent, respectively. The wedge between the top and bottom quartiles grows with the return horizon, to 9.41 and 4.56 percent at the top and bottom quartiles of the ten-year return horizon. Note that when keeping the (non-annualized) return constant, the probability of sale in the next year decreases with the return horizon. However, expected (non-annualized) returns increase with the horizon, so that in expectation, paintings are more likely to sell in the next year as the horizon grows.

The fourth result revealed by Figure 2 is a further asymmetry between gains and losses: the sale probability function is steeper and mildly concave for losses, and flatter and convex for gains, although these effects are rather mild for most of the empirically relevant range of returns. Appendix D shows that the general shape of the sale probability function is the same in the early and late part of the sample.

Model B expands the set of variables in the selection equation with three more variables: the size of the painting, an indicator whether the artist deceased in the past two years, and the

[17] If there are (unobserved) private sales in between two auctions, this may affect the estimated sales probability function. The direction of change is not quite clear, though. For example, when we observe a 30% return at a 10-year horizon, with an (unobserved) intermediate sale this could in fact be, say, two 15% returns at 5-year horizons. This means that sale probabilities may be more sensitive to low to moderate returns, implying that the sample selection problem may be even stronger, particularly since these are the most common return observations. At the same time, the selection effect may be attenuated for higher returns.

growth in worldwide GDP. Table 3 shows that larger paintings are more likely to sell than small ones, because the larger pieces in our data tend to be more ideal sizes to display. Paintings are not significantly more likely to sell within two years of the artist's death, despite popular belief that the passing away of an artist temporarily raises visibility of his or her artwork. This may be because most artists's deaths do not come as a surprise to the market. High worldwide GDP growth is associated with higher sales probability, as more wealth translates into money flowing into art, consistent with Goetzmann (1993), Hiraki et al. (2009), and Goetzmann et al. (2011). Most importantly, though, Table 3 shows that the relation between r_i^p and sales is robust to the inclusion of these additional variables. The sale probabilities for Model B and the other model specifications that we consider below look very similar to Model A, so we do not report them separately.

In Model C we consider whether there are differences between art styles. Buelens and Ginsburgh (1993) find that different styles are in favor in different periods, and Penasse, Renneboog, and Spaenjers (2014) show that the correlation between prices and sales volume of paintings may be driven by a fad component. Styles that are "hot" may be subject to more speculative buying and selling, resulting in more realizations of short-term gains as predicted by the model of Lovo and Spaenjers (2014). In addition, auction houses' incentives are to maximize revenue, and they organize sales around a particular style they believe will generate high revenue. To assess whether more sales of a style increase the sale probability of any single piece of that style, we include the relative market share of a style in the full BASI data set, defined as the present year's share of total sales volume divided by its five-year historical average share (to control for a baseline level of sales). A style that is currently popular has a high relative market share.[18] The average (median) relative share is 1.03 (1.01), with a standard deviation of 0.21. We also include interactions of the relative market share with r_i^p, as styles that peaked in past years may be more subject to the sample selection bias.

The insignificant coefficient on relative market share in Table 3 shows that a style's current popularity does not uniformly raise the probability of a sale of a painting of that style. For paintings that have dropped in value since purchase, there is also no statistical difference in sale probabilities between popular and unpopular styles, potentially because liquidity shocks, taxes or other non-demand related factors dominate for losses. Rather, the effects of popularity are concentrated in the paintings that have risen in value. The discontinuity at $r^p = 0$ is larger for popular styles: at the one year horizon the jump in probability is 1.09 percentage points for styles that are one standard deviation above average relative share, compared to 0.82 percentage points for styles that are one standard deviation below average. For popular styles, small gains - up to returns of 25 percent - are more likely to realize. This could be driven by speculative selling, conform Lovo and Spaenjers (2014). For gains over 25 percent, the relation switches and unpopular styles are more likely to sell.[19] In other words, people tend to wait for higher gains before they sell unpopular styles. Note that since larger gains typically only occur when smaller gains compound, the popular style paintings are overall more likely to sell.[20]

This result helps explain the high correlations between styles' market shares and returns in Table 2. Styles are more prone to sample selection biases when they are popular. Not only do higher returns raise the sale probability (as they do in Models A and B), the interaction with

[18] Our results are robust to using the relative dollar value of sales (rather than the number of paintings sold) as the measure of popularity.
[19] The switching point of 25 percent is consistent across horizons.
[20] High sales volume may also proxy for adverse selection or liquidity, but these explanations appear inconsistent with the results. If high volume proxies for low adverse selection, then styles with low adverse selection experience higher returns. If volume proxies for liquidity, then there is an illiquidity discount rather than a premium.

relative market share amplifies this effect when the style is in swing, as happened for example for Impressionist paintings in the early 1970s and late 1980s, and for Post-war and Contemporary paintings in the late 1980s and the mid-2000s. Old Masters, and especially "Other" style paintings and top-selling artists do not experience periods of high market shares (relative to their baselines) during the sample period. The relation between returns and sales volume is less pronounced for these categories, hence their lower correlations in Table 2.

3.3. Art indices

Figure 3 Panel A plots two estimated art price indices that ignore selection bias. The first index is constructed from the standard repeat sales regression estimated by GLS, weighing each observation by the inverse of the square root of the time between sales, to correct for potential heteroskedasticity (the "GLS index"). The second is a MCMC specification that ignores selection by forcing all $\alpha_1, \ldots, \alpha_K$ in equation (3) to equal zero (the "MCMC index"). We normalize the index value to 100 in the year 1960. The GLS and MCMC indices practically coincide, mitigating concerns about distributional assumptions of the MCMC estimator.

[Please insert Figure 3 here]

Next, we use Models A, B, and C of Table 3 to construct three selection-corrected indices, which we denote Indices A, B, and C, respectively. Panel A of Figure 3 also plots the time-series of the selection-corrected price indices. The selection correction is quite robust across models: the differences among the three selection-corrected indices are generally less than ten index points. Compared to the non-corrected indices, however, the differences are striking. The selection-corrected indices are markedly lower than the non-corrected indices, consistent with the earlier finding that gains are overall more likely to realize than losses. The peak in the indices in 1990, which occurs at an index level of 1,775 in the non-selection-corrected model, occurs at around 860 in the corrected indices. Following the Japanese real estate collapse in the early 1990s, the non-corrected indices bottom out at 1,275 in 1993, while the selection-corrected ones are around 660. The 2007 peak in the non-corrected models is near 5,000, versus around 1,625 in the selection-corrected models. By 2013, the end of our sample period, the price indices have recovered from the dip in the global financial crisis of 2008 to 2009, and the non-corrected index is around 5,429, compared to around 1,895 for the selection-corrected models.

The difference between the selection-corrected and non-corrected indices grows fastest when art prices are rising, i.e., in the years up to 1990, the boom from 1993 to 2000, and from 2001 up to the financial crisis in 2008. This happens because in "normal" times when the market is rising, gains are more likely to realize, and the non-corrected index overstates the price rise. Conversely, during the few periods when art values experience large drops, such as the bursting of the Japanese bubble in the early 1990s, the market crash of 2000, and the financial crisis of 2008 and 2009, the difference between the indices narrows somewhat. When there are few gains and large losses, those losses are relatively more likely to realize due to the V-shaped probability function. The selection bias thus reverses, and the non-corrected index overstates the decline, pulling the index closer to the selection-corrected one. This also suggests that the non-corrected index is excessively volatile, and we show below that its return volatility is indeed higher than the selection-corrected index.

Next, Figure 4 shows the selection-corrected indices estimated by style. Styles mostly follow a common pattern, peaking in 1990 and again in 2008, but there are some differences. Post-war and Contemporary paintings did well in the booms of the late 1980s and especially the early and mid-2000s. American and 19th Century European paintings also performed well, and at a steadier, less volatile pace. Impressionist and Modern paintings show large increases in the 1980s but are hit heavily in 1990, in line with the popular interpretation of art observers that the strong yen in the 1980s and the bursting of the Japanese real estate bubble in 1990 had strong effects on the prices of Impressionist paintings (see, for example, Wood, 1992). Old Masters did not increase much in value over the sample period, compared to the other styles.

[Please insert Figure 4 here]

The most striking result is the performance of the top 100 artists portfolio (as defined by sales over the prior decade), which outperformed all styles, and did well even after the financial crisis, when style indices remained mostly flat. 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, defined as paintings in the right tail of the price distribution, and generally find that they underperform (Pesando, 1993, Mei and Moses, 2002, 2005), or find mixed effects (Ashenfelter and Graddy, 2003, Pesando and Shum, 2008). Renneboog and Spaenjers (2013) criticize this approach as identifying the most overpriced paintings rather than true masterpieces. They instead consider artists that are most frequently mentioned in five editions of Gardner's art history textbooks, and find evidence that these master artists outperform. Our definition of top-selling artists most resembles theirs, and we show below that, despite the more volatile returns to this strategy, this result holds even after controlling for selection bias.

3.4. Buy-ins

Including the additional information in buy-ins yields parameter estimates that are close to the sales-only sample estimates in Table 3. For brevity we report them in Appendix E. Although the sale probability function is nearly identical to the one shown in Figure 2, the inclusion of buy-in information does have an impact on the estimated indices. Panel B of Figure 3 compares the Model C index estimated with and without the buy-in information. We renormalize the indices to 100 in 2007, the first year in which we have buy-in information, and focus on the 2007 to 2013 period. The buy-in index gradually diverges from the sales-only index up until 2012, at which point it has dropped 3.3 index points (or 3.1%) below the sales-only index. The two indexes converge in 2013, however, and the difference at the end of the sample period is therefore rather small. With a longer time-series of buy-ins it is likely that the differences will be larger.

4. Optimal portfolio allocation

In this section, we show the importance of sample selection for performance evaluation and optimal portfolio allocation. For brevity, we focus on the results from the sales-only sample (i.e., without the buy-in information), leaving robustness checks to Appendix F.

[Please insert Table 4 here]

Table 4 reports means, standard deviations and Sharpe ratios of the annual arithmetic returns on the art indices. The standard repeat sales GLS index has an average annual return of 8.7 percent with a standard deviation of 13.8 percent, and an annual Sharpe ratio of 0.268. The non-selection-corrected MCMC index returns are nearly identical. In contrast, the selection-corrected index A has an average return of 6.3 percent, which is 28 percent lower than the non-corrected index. The standard deviation is 11.4 percent. The material excess volatility of the non-corrected index is due to the time-variation in the strength of sample selection, as argued above.

The Sharpe ratio of the selection-corrected index is 0.111, or 59 percent lower than the non-corrected index. The return properties of index B and C are nearly identical to index A.[21]

The selection-corrected returns are more representative of the experience of an investor who has invested in a well-diversified (non-targeted) portfolio of paintings, since the standard non-selection-corrected index implicitly assumes that an investor can either pick "winners" that rise in value and are thus more likely to sell, or that there is no selection problem and that the investor's holdings follow the same price path as the paintings that come to auction.[22] The selection-corrected indices do not make such assumptions, but rather measure the rise in value of a representative portfolio of paintings, both those that sold and those that did not.[23]

Table 4 also shows descriptive statistics of returns of a broad portfolio of global equities (the total return on global equities from Dimson, Marsh, and Staunton, 2002)[24], corporate bonds (the Dow Jones corporate bond return index), commodities (the World Bank Global Economic Monitor commodities index), and real estate (the U.S. residential real estate index from Shiller, 2009) over the sample period. For the risk-free asset we use the global Treasury bill rate at the beginning of the year from Dimson, Marsh, and Staunton. Being our risk-free asset, we do not report its Sharpe ratio.

Despite the low Sharpe ratios of the selection-corrected art indices relative to stocks, corporate bonds, and commodities (of 0.36, 0.42, and 0.16, respectively), investing in paintings may still be useful for constructing optimal portfolios if the correlations between art and the other asset classes are low. Table 4 shows that the correlation between art and the other assets are

[21] Appendix G shows that the non-linearities in the sale probability function are economically important for art returns, compared to a linear probability function.

[22] Based on the selection model results, it takes about 30 to 40 paintings to exhaust the diversification benefits in an equally weighted portfolio, but a portfolio of as few as 15 paintings could be considered well diversified.

[23] One caveat is that our data include an unidentified number of estate, divorce, and bankruptcy sales (the "3D's": debt, death, and divorce). If such sales fetch lower prices than regular sales, then the returns from the selection-corrected model are representative for investors only insofar that they are subject to these events at similar rates.

[24] We are grateful to Elroy Dimson, Paul Marsh, and Mike Staunton for generously sharing their data with us.

less than 0.3 for almost all art indices, and statistically no different from zero except for the correlation with real estate.

We construct optimal portfolios using common base case assumptions from the literature. First, investors have mean-variance utility. Second, borrowing and short sales are not allowed. Third, there are no transaction costs to constructing the portfolios. Fourth, there is no illiquidity premium on paintings. Fifth, investing in the art index does not provide the investor with access to the artworks underlying the index, and thus there is no consumption utility of owning art. We use Dimson's (1979) method with one year leads and lags of returns to correct for first-order autocorrelation due to time-aggregation of sales (Working, 1960, and Schwert, 1990).

[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). An investor who does not correct for selection bias in art returns would want to assign 31 percent of the portfolio's weight to art. The perceived portfolio Sharpe ratio of 0.64 is 10 percent higher than the Sharpe ratio of 0.59 that is achieved with stocks alone, an economically significant improvement.

In contrast, an investor who corrects for sample selection optimally assigns zero weight to paintings across all selection models A, B, and C. This stark result underscores the importance of correcting for sample selection when making optimal portfolio decisions. Had the investor followed GLS allocations, with 31% of the portfolio in art, the realized portfolio Sharpe ratio based on the selection-corrected art returns would be 0.50. Compared to the 0.59 Sharpe ratio on a portfolio excluding art, this is a loss of 14% (this result is not tabulated).

Panels B and C of Table 5 show the portfolio allocations for a mean-variance utility investor with a risk aversion coefficient equal to two or ten, respectively. The results are consistent with the tangency portfolio results, though somewhat weaker for the case of risk aversion equal to two.

One issue with the above analysis is that it may not be feasible to invest in such an art index, since one cannot purchase a fraction of every painting, unlike stocks or bonds. At best an investor can build a portfolio that contains a broad set of paintings that mimics the makeup of the art market. Such a portfolio must have a lower Sharpe ratio than the index. Other frictions such as transaction costs and illiquidity will further depress attainable Sharpe ratios, strengthening the conclusion that investing in a broad portfolio of paintings is not economically profitable.

Despite the above result about investing in the broad art market, it is perhaps economically sensible to construct a portfolio that focuses on particular styles of paintings, or on top-selling artists. Indeed, Table 5 shows that the American style and the Top 100 artists receive positive portfolio weights after adjusting for selection bias. The Post-war and Contemporary style is not in the optimal portfolio, despite having a higher Sharpe ratio than American paintings, because its correlations with the other asset classes are relatively high. The Sharpe ratio of the tangency portfolio improves by nearly 9 percent to just shy of 0.64, relative to the 0.59 Sharpe ratio when art is excluded. Thus, there appears to be improvement in performance when considering narrower, targeted categories of paintings. Still, we should be cautious in drawing strong conclusions from this exercise, as the styles' relative performance is not consistent over time (as seen in Figure 4), ignores transaction costs, and is based on only five decades of data.

Appendix F shows that the portfolio results are robust to a range of alternative assumptions and measures, such as adding the buy-in information, adjusting for illiquidity using the model by

Ang, Papanikolaou, and Westerfield (2014), unsmoothing the index returns instead of using the Dimson correction (e.g., Campbell, 2008, and Renneboog and Spaenjers, 2013), and assuming power utility rather than mean-variance utility.

5. Conclusion

We estimate an empirical model that adjusts for selection bias in illiquid asset markets with endogenous sales, using a large data set of art auctions. Our model generalizes the prior literature on selection bias in important ways. Most importantly, we allow for a non-linear relation between returns and sale probabilities, and we use information from buy-ins, where paintings came up for auction but did not sell. We find that the increased flexibility is economically important. The relation between art returns and sale probabilities is V-shaped, i.e., large gains and losses are more likely to realize than small gains and losses. The relation is asymmetric, and in normal times gains are over-represented in the data, largely due to an upward jump in sale probabilities as marginal losses become marginal gains. This results in a disposition effect as hypothesized by Goetzmann (1993, 1996). We also show that there is time and cross-sectional variation in the strength of the sample selection bias. During significant downturns in the art market, selection bias may even switch sign, as large losses become over-represented in the data. In addition, styles are more prone to selection bias when they are popular, consistent with an increase in speculative buying and selling behavior.

Sample selection bias has a first-order impact on art indices, lowering the average annual return by 28 percent, from 8.7 percent for a standard repeat sales index to 6.3 percent for selection-corrected indices. The standard deviation of the non-corrected index returns is 2.4 percentage points (i.e., 21 percent) higher than the corrected index. This excess volatility is due to the time-variation in the strength of selection bias. The annual Sharpe ratio of the corrected index drops nearly 60 percent, from 0.27 to 0.11. The implications are that an investor would not find it attractive to invest in a portfolio that is representative for the broad art market, unless she derives substantial non-monetary utility from owning and enjoying art. Had she ignored selection bias, she would have allocated a non-negligible share of her portfolio to art. We find some suggestive evidence that a strategy targeted at certain styles or at top-selling artists may be optimally included in an investment portfolio of art, stocks, bonds, real estate, and commodities.

More broadly, this paper highlights the importance of accounting for sample selection for performance evaluation and portfolio optimization of illiquid assets with endogenous sales. Whether the selection correction is quantitatively as large in other asset classes (e.g., real estate, private equity, and corporate bonds) as in art, is an important empirical question that we leave for future work. The methodology developed in this paper should prove helpful in answering these questions, as it naturally applies to these other settings.

References

Adam, G., 2008. When Brueghel met Schnabel. *Financial Times*, April 21, 2008.

Ang, A., Papanikolaou, D., and Westerfield, M., 2014. Portfolio choice with illiquid assets. *Management Science* 60, 2737-2761.

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.

Barberis, N., 2013. Thirty years of prospect theory in economics: A review and assessment. *Journal of Economic Perspectives* 27, 173-196.

Barberis, N., and Xiong, W., 2012. Realization utility. *Journal of Financial Economics* 104, 251-271.

Baumol, W., 1986. Unnatural value: Or art investment as floating crap game. *American Economic Review* 76, 10-14.

Beggs, A., and Graddy, K., 2008. Failure to meet the reserve price: The impact on returns to art. *Journal of Cultural Economics* 32, 301-320.

Ben-David, I., and Hirschleifer, D., 2012. Are investors really reluctant to realize their losses? Trading responses to past returns and the disposition effect. *Review of Financial Studies* 25, 2485-2532.

Berkson, J., 1950. Are there two regressions? *Journal of the American Statistical Association* 45, 164-180.

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

Burton, B.T., and Jacobsen, J.P., 1999. Measuring returns on investments in collectibles. *Journal of Economic Perspectives* 13, 193-212.

Case, K.E., Pollakowski, H.O., and Wachter, S.M., 1991. On choosing among house price index methodologies. *Real Estate Economics* 19, 286-307.

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.

Campbell, J., Giglio, S., and Pathak, P., 2011. Forced sales and house prices. *American Economic Review* 101, 2108-2131.

Cochrane, J., 2005. The risk and return of venture capital. *Journal of Financial Economics* 75, 3-52.

Collins, A., Scorcu, A.E., and Zanola, R., 2007. Sample selection bias and time instability of hedonic art price indices. *Working paper No. 610*, University of Bologna.

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

Dimson, E., Marsh, P., and Staunton, M., 2002. Triumph of the optimists: 101 years of global investment returns. 1st ed. Princeton, New Jersey: Princeton University Press.

Frazzini, A., 2006. The disposition effect and under-reaction to news. *Journal of Finance* 61, 2017-2046.

Frey, B.S., and Eichenberger, R., 1995. On the rate of return in the art market: Survey and evaluation. *European Economic Review* 39, 528-537.

Gatzlaff, D.H., and Haurin, D.R., 1997. Sample selection bias and repeat-sales index estimates. *Journal of Real Estate Finance and Economics* 14, 33-50.

Gatzlaff, D.H., and Haurin, D.R., 1998. Sample selection and biases in local house value indices, *Journal of Urban Economics* 43, 199-222.

Genesove, D., and Mayer, C., 2001. Loss aversion and seller behavior: Evidence from the housing market. *Quarterly Journal of Economics* 116, 1233-1260.

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

Goetzmann, W., and Peng, L., 2006. Estimating house price indexes in the presence of seller reservation prices. *Review of Economics and Statistics* 88, 100-112.

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

Grinblatt, M., and Keloharju, M., 2001. What makes investors trade? *Journal of Finance* 56, 589–616.

Grossman, S.J., and Stiglitz, J.E., 1980. On the impossibility of informationally efficient markets. *American Economic Review* 70, 393-408.

Hartzmark, S.M., and Solomon, D.H., 2012. Efficiency and the disposition effect in NFL prediction markets. *Quarterly Journal of Finance* 2, 1250013.

Haurin, D.R., and Hendershott, P.H., 1991. House price indexes: Issues and results. *Real Estate Economics* 19, 259–269.

Heath, C., Huddart, S., and Lang, M., 1999. Psychological factors and stock option exercise. *Quarterly Journal of Economics* 114, 601-627.

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

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

Hiraki, T., Ito, A., Spieth, D.A., and Takezawa, N., 2009. How did Japanese investments influence international art prices? *Journal of Financial and Quantitative Analysis* 44, 1489-1514.

Hwang, M., and Quigley, J.M., 2004. Selectivity, quality adjustment and mean reversion in the measurement of house values. *Journal of Real Estate Finance and Economics* 28, 161-178.

Hwang, M., Quigley, J.M., and Woodward, S.E., 2006. An index for venture capital, 1987-2003. *The B.E. Journals in Economic Analysis & Policy* 4, 1-43.

Ingersoll, J., and Jin, L., 2013. Realization utility with reference-dependent preferences. *Review of Financial Studies* 26, 723-767.

Jacquier, E., Polson, N.G.P., and Rossi, P.E., 1994. Bayesian analysis of stochastic volatility models. *Journal of Business and Economic Statistics* 94, 371–417.

Jud, G.D., and Seaks, T.G., 1994. Sample selection bias in estimating housing sales prices. *Journal of Real Estate Research* 9, 289-298.

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., 2014. Estimating loan-to-value distributions. *Working paper*, Stanford University and Columbia University.

Lovo, S., and Spaenjers, C., 2014. Unique durable assets. *Working Paper*, HEC Paris.

Mandel, B.R., 2009. Art as an investment and conspicuous consumption good. *American Economic Review* 99, 1653-1663.

McAndrew, C., 2014. TEFAF art market report 2014. Arts Economics, Dublin (Ireland).

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

Mei, J., and Moses, M., 2005. Vested interest and biased price estimates: Evidence from an auction market. *Journal of Finance* 60, 2409-2435.

Meng, J., 2014. Can prospect theory explain the disposition effect? A new perspective on reference points. *Working paper*, Peking University.

Munneke, H.J., and Slade, B.A., 2000. An empirical study of sample-selection bias in indices of commercial real estate. *Journal of Real Estate Finance and Economics* 21, 45-64.

Munneke, H.J., and Slade, B.A., 2001. A metropolitan transaction-based commercial price index: A time-varying parameter approach. *Real Estate Economics* 29, 55-84.

Odean, T., 1998. Are investors reluctant to realize their losses? *Journal of Finance* 53, 1775-1798.

Pearce, B., 2014. Adapting and evolving: Global venture capital insights and trends 2014. Ernst & Young.

Pénasse, J., Renneboog, L., and Spaenjers, C., 2014. Sentiment and art prices. *Economics Letters* 122, 432-434.

Peng, L., 2001. Building a venture capital index. *Working paper*, University of Colorado.

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

Pesando, J., and Shum, P., 2008. The auction market for modern prints: Confirmation, contradictions, and new puzzles. *Economic Inquiry* 46, 149-159.

Picinati di Torcello, A., and Petterson, A., 2014. Art & Finance Report 2014, Deloitte Luxembourg and ArtTactic.

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

Schwert, G.W., 1990. Indexes of U.S. stock prices from 1802 to 1987. *Journal of Business* 63, 399-442.

Shefrin, H., and Statman, M., 1985. The disposition to sell winners too early and ride losers too long: Theory and evidence. *Journal of Finance* 40, 777-791.

Shiller, R.J., 2009. Irrational exuberance. 2nd ed. Princeton, NJ: Princeton University Press.

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.

The Economist, 2013. Freeports: Über-warehouses for the ultra-rich. November 23, 2013.

Wilkoff, S., 2013. A municipal bond market index based on a repeat sales methodology. *Working paper*, U.S. Securities and Exchange Commission.

Wood, C., 1992. The bubble economy. The Japanese economic collapse. London: Sidgwick & Jackson.

Woodward, S.E., and Hall, R.E., 2003. Benchmarking the returns to venture. *Working paper*, Sand Hill Econometrics and Stanford University.

Working, H., 1960. Note on the correlation of first differences of averages in a random chain. *Econometrica* 28, 916-918.

Zanola, R., 2007. The dynamics of art prices: The selection corrected repeat-sales index. *POLIS working paper No. 85*, University of Eastern Piedmont.

Figure 1: Number of sales

This figure shows the number of auction sales of paintings in the repeat sales sample, by calendar year. We distinguish between the first, second, and third or more sales of an artwork.

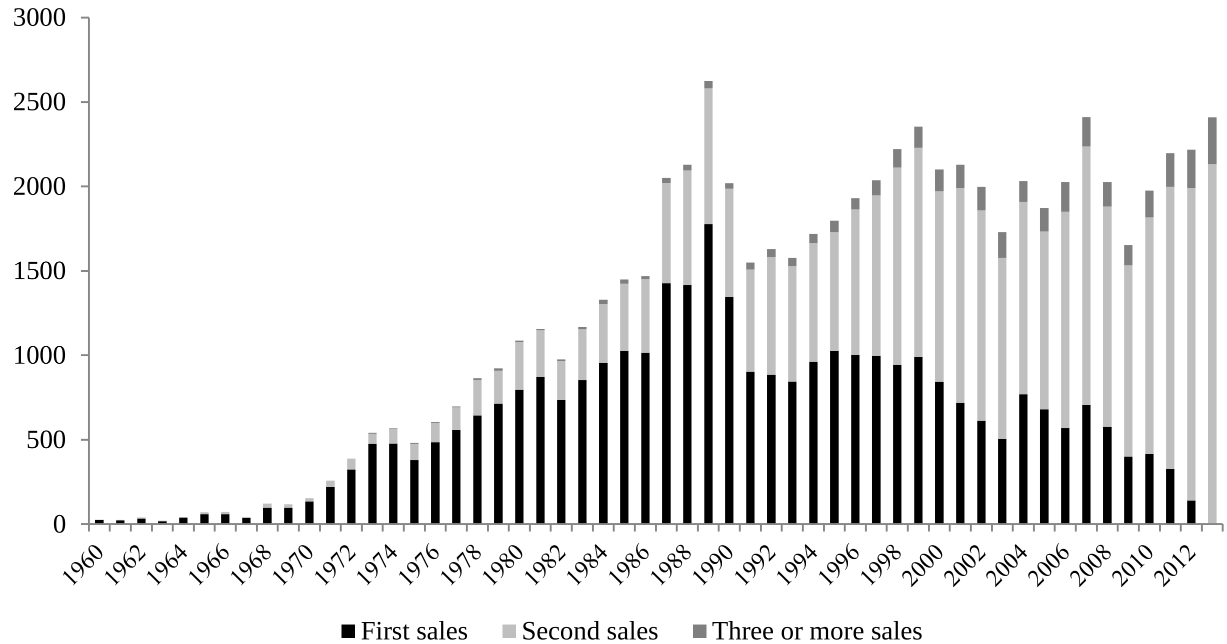

Figure 2: Probability of sale

This figure shows the model-implied probability (in percent, on the vertical axis) that a painting will sell in the next year, measured one year (left plot), five years (middle plot), or ten years (right plot) after a prior sale, as a function of the return since the prior sale (non-annualized, in percent, on the horizontal axis). The sale probabilities are computed for the selection model specification A in Table 3, not using the buy-in information. *E(r)* is the estimated mean of the return distribution for the relevant horizon, based on the estimated price paths of all paintings, and is marked with an 'x' on the horizontal axis of each plot. Similarly, *p25* and *p75* are the 25th and 75th percentiles of the return distribution, and are also marked with 'x's.

Figure 3: Art price indices

This figure shows repeat sales arithmetic art price indices, normalized to an index value of 100 in 1960. Panel A shows two indices not corrected for sample selection (*GLS* and *MCMC*), and three indices that are corrected for sample selection (*A*, *B*, and *C*). The *GLS* index is the standard repeat sales regression index 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., all α_k for $k = 1 \dots K$ in equation (3) are set to zero. Models A through C correct for sample selection, and correspond to the specifications of the selection equation as shown in Table 2. Panel B shows two estimated of Model C in Table 3, normalized to 100 in 2007, where the solid line is estimated on the sample that excludes buy-ins, while the striped line uses the buy-in information.

Panel A. Non-selection-corrected and selection-corrected price indices (sales only)

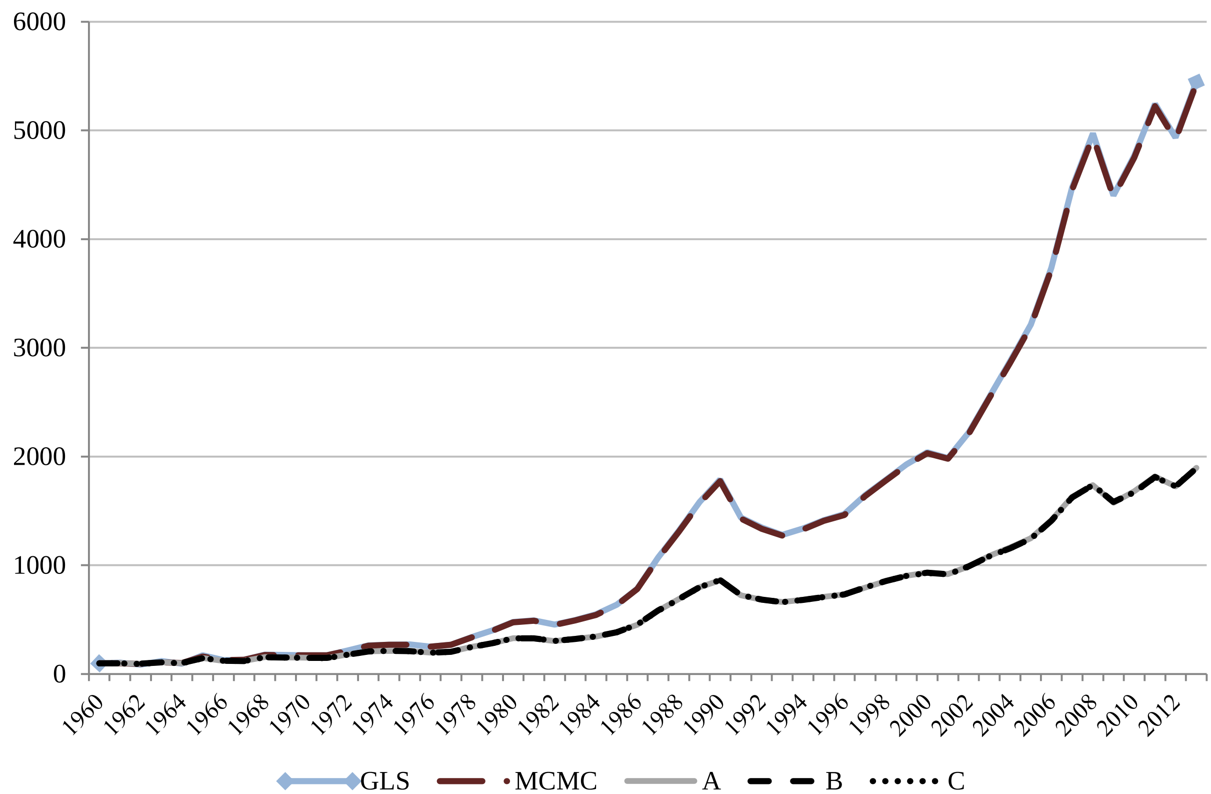

Panel B. Selection-corrected price indices (sales and buy-ins)

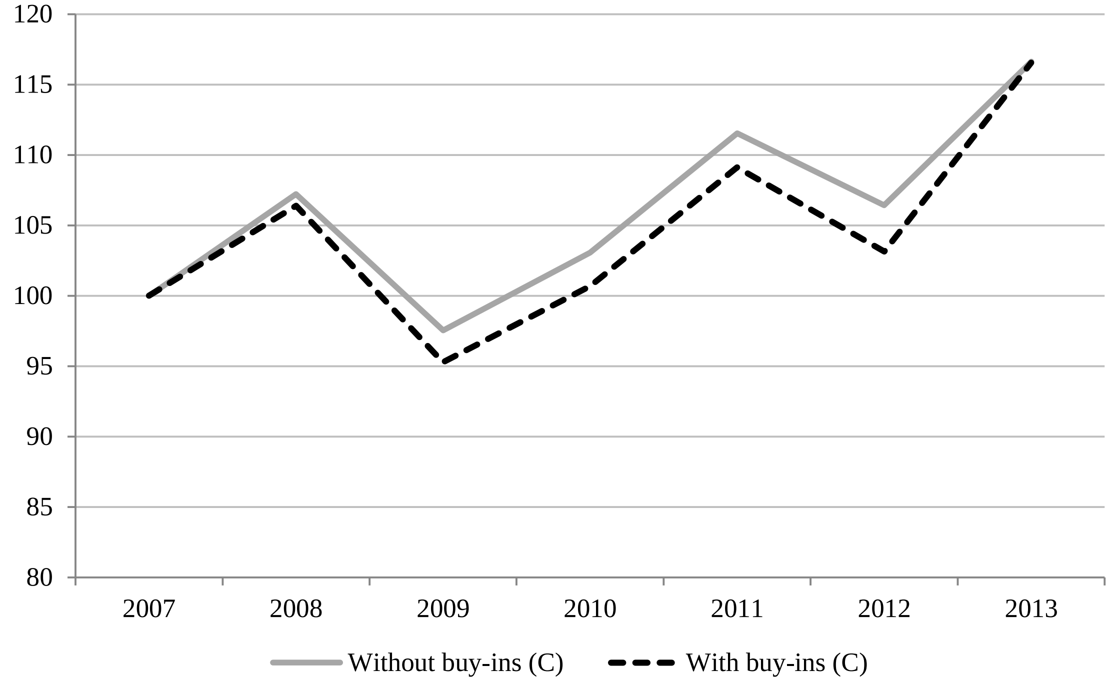

Figure 4: Selection-corrected price indices per style

This figure shows selection-corrected price indices for each style classification, normalized to an index value of 100 in 1960. *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 decade prior to the year of sale.

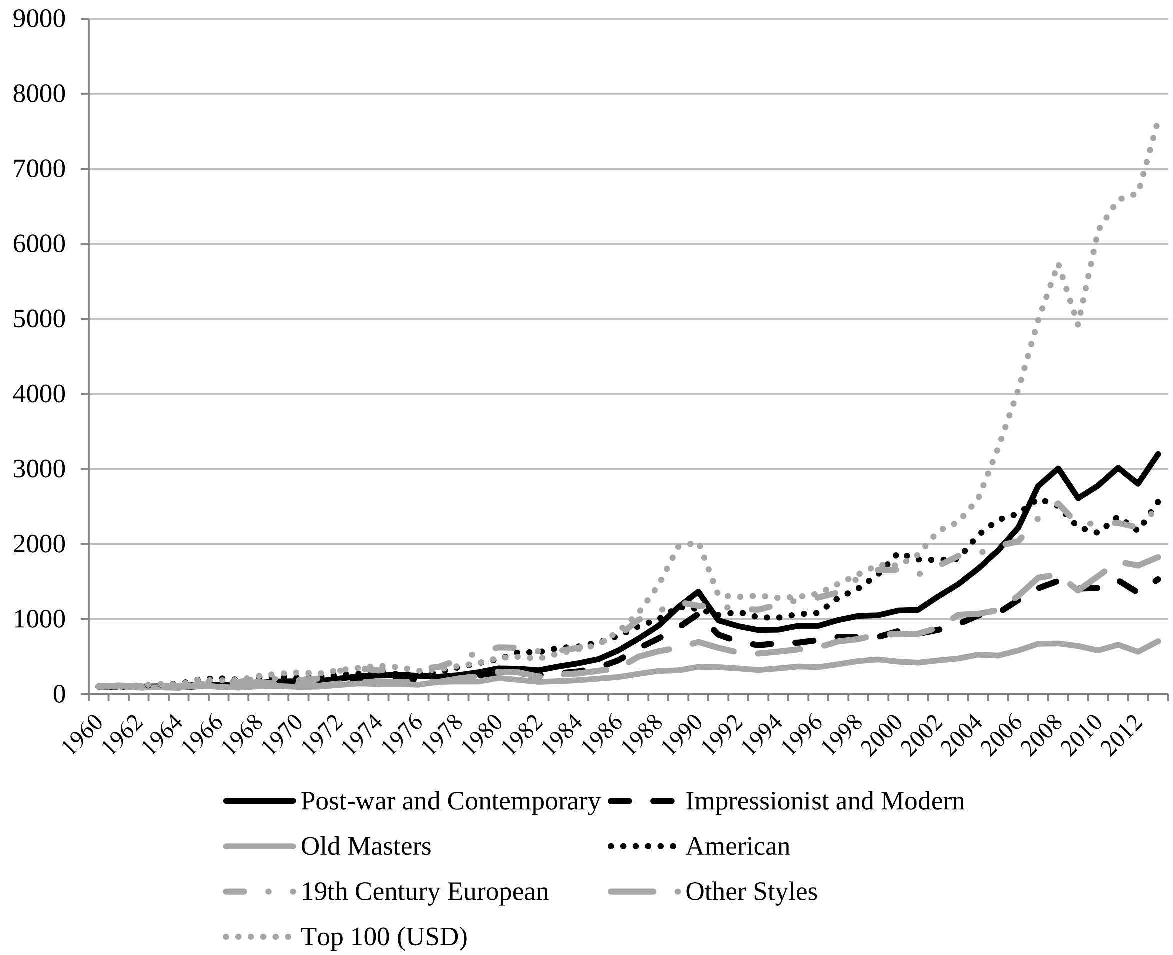

Table 1: Summary statistics

This table reports summary statistics for the repeat sales sample of paintings in the Blouin Art Sales Index (BASI) dataset from 1960 to 2013. Panel A presents descriptive statistics of the 32,928 paintings that sold at least twice during the sample period. The unit of observation is a sale of a painting at auction. *Hammer price* is the auction price in thousands of U.S. dollars. *Low Estimate* and *High Estimate* are the auction house's low and high price estimates, respectively, as a percentage of the hammer price. *Surface* is the surface of the painting in thousands of squared millimeters. *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. *Relative share* is the market share (in terms of sales) of the painting's style in the year of sale compared to the style's average market share in the five years prior to the sale. *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 the decade prior to the year of sale, and zero otherwise. The remaining variables in Panel A represent style classifications. Panel B shows descriptive statistics of the sale-to-sale returns in the data.

Panel A. Descriptive statistics (69,103 sales)

	Mean	Median	St. Dev.
Hammer price ($000s)	150.6	14.5	923.8
Low estimate (% of hammer price)	85.64%	81.63%	76.50%
High estimate (% of hammer price)	118.83%	114.29%	83.23%
Surface	630.3	352.3	1,161.1
Deceased < 2yrs	1.72%		
Christie's	32.58%		
Sotheby's	32.67%		
London	26.87%		
New York	29.11%		
Post-war and Contemporary	14.32%		
Impressionist and Modern	25.80%		
Old Masters	11.23%		
American	10.23%		
19^{th} Century European	21.20%		
Other Styles	17.22%		
Top 100 Artists	15.20%		

Panel B. Sale-to-sale returns (36,175 returns)

	Mean	Median	St. Dev.
Sale-to-sale return			
Arithmetic return	109.37%	47.06%	187.38%
Log return	45.93%	38.57%	71.99%
Years between sales	9.42	7.33	7.46
Number of sales per painting	2.10		
Annualized sale-to-sale return			
Arithmetic return	12.37%	6.66%	21.85%
Log return	5.59%	4.80%	13.49%

Table 2: Relation between annualized returns and market shares by style

Panel A shows the correlation between the proportion of sales that occur in a given year and the mean (left column) and median (right column) annualized log sale-to-sale return, computed over the returns for which the second sale falls in the same year. The proportion of sales is calculated as the number of sales in the year as a percentage of all sales over the 1960 to 2013 sample period. Panel B shows the correlation between the yearly market share of each style and the mean (left column) and median (right column) annualized log sale-to-sale return for the same style. The market share of a style is the total sales of that style in a given year relative to all sales in the same year, calculated from the full BASI dataset. Panel C shows the coefficients of a regression of the yearly style market shares on the annualized sale-to-sale returns by style, and style fixed effects. The sample period for Panels B and C is from 1972 to 2013 to allow for sufficient second sales to calculate meaningful mean and median sale-to-sale returns. Standard errors are in parentheses. ***, ** and * indicate statistical significance at the 1, 5 and 10% level, respectively.

Panel A. Correlation coefficient between proportion of sales and return

	Mean annualized sale-to-sale return	Median annualized sale-to-sale return
All paintings	0.409	0.252

Panel B. Correlation coefficients between market share and return, by style

	Mean annualized sale-to-sale return	Median annualized sale-to-sale return
Post-war and Contemporary	0.182	0.237
Impressionist and Modern	0.429	0.477
Old Masters	0.202	0.126
American	0.386	0.482
19th Century European	0.471	0.503
Other Styles	0.086	-0.063
Top 100 Artists	0.067	0.043

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

	I	II
Mean annualized sale-to-sale return	0.163***	
	(0.048)	
Median annualized sale-to-sale return		0.185***
		(0.056)
Style fixed effects	Yes	Yes
Adjusted R^2	77.5%	77.4%
Number of observations	293	293

Table 3: Selection equation coefficients

This table presents the parameter estimates of three specifications of the selection equation (equation (3) in the text). *Return* is the natural logarithm of the return since the prior sale of a painting (non-annualized). *Relative share* is the market share (in terms of sales) of the painting's style in the year of sale compared to the style's average market share in the five years prior to the sale. *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 the Historical Statistics of the World Economy. The other variables are as defined in Table 1. *Sigma* is the standard deviation of the idiosyncratic error term in equation (1). Standard errors are in parentheses. ***, ** and * indicate statistical significance at the 1, 5 and 10% level, respectively.

	A		B		C	
Return > 0	0.124	***	0.122	***	0.047	
	(0.013)		(0.013)		(0.048)	
(Return > 0)					0.075	*
* relative share					(0.046)	
(Return<0) * return	-0.683	***	-0.691	***	-0.650	***
	(0.049)		(0.049)		(0.152)	
(Return<0) * return					-0.041	
* relative share					(0.141)	
(Return<0) * return^2	-0.383	***	-0.382	***	-0.406	**
	(0.038)		(0.038)		(0.160)	
(Return<0) * return^2					0.023	
* relative share					(0.152)	
(Return>0) * return	0.149	***	0.151	***	0.496	***
	(0.024)		(0.024)		(0.136)	
(Return>0) * return					-0.342	**
* relative share					(0.133)	
(Return>0) * return^2	0.231	***	0.235	***	0.068	
	(0.015)		(0.015)		(0.077)	
(Return>0) * return^2					0.163	**
* relative share					(0.075)	
Relative share					0.001	
					(0.020)	
Time (years)	-0.014	***	-0.015	***	-0.015	***
	(0.001)		(0.001)		(0.001)	
Time squared	0.000	***	0.000	***	0.000	***
	(0.000)		(0.000)		(0.000)	
Log (surface)			0.010	***	0.010	***
			(0.003)		(0.003)	
Deceased < 2 years			0.041		0.041	
			(0.027)		(0.027)	

World GDP growth			2.047	***	2.052	***
			(0.183)		(0.182)	
Style fixed effects	Yes		Yes		Yes	
Sigma	0.199	***	0.199	***	0.199	***
	(0.001)		(0.001)		(0.001)	

Table 4: Descriptive statistics of annual index returns

This table reports descriptive statistics of the annual arithmetic returns to indices of paintings and other assets over the period 1961 to 2013. *GLS* is the standard repeat sales 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 3 (without the buy-in information). *Stocks* is the global equity total return from Dimson, Marsh, and Staunton (2002). *Corporate bonds* is the Dow Jones corporate bond return index. *Commodities* is the return on the World Bank GEM commodities index. *Real estate* is returns on the U.S. residential real estate index from Shiller (2009). *Treasuries* are global Treasury bill rates from Dimson, Marsh, and Staunton (2002). Sharpe ratios are annualized. ***, ** and * indicate statistical significance of the correlation coefficients at the 1, 5 and 10% level, respectively.

	Mean	St. dev.	Sharpe ratio	Excess return correlation with			
				Stocks	Corp. bonds	Commo-dities	Real estate
Returns on non-selection-corrected art indices:							
GLS	8.71%	13.76%	0.268	0.192	-0.186	0.106	0.227
MCMC	8.68%	13.92%	0.263	0.198	-0.208	0.136	0.292**
Returns on selection-corrected art indices:							
A	6.29%	11.42%	0.111	0.222	-0.199	0.132	0.286**
B	6.28%	11.34%	0.111	0.216	-0.198	0.131	0.283**
C	6.28%	11.35%	0.111	0.220	-0.197	0.132	0.283**
Returns on selection-corrected art sub-indices							
Post-war and Contemporary	7.43%	11.63%	0.208	0.136	-0.172	0.151	0.347**
Impressionist and Modern	6.09%	13.30%	0.080	0.103	-0.230*	0.199	0.222
Old Masters	4.56%	13.75%	-0.033	0.222	-0.107	0.042	0.195
American	6.83%	10.28%	0.176	0.076	-0.298**	0.102	0.234*
19th Century European	6.81%	11.70%	0.153	0.304**	-0.155	0.054	0.233*
Other Styles	6.53%	13.92%	0.109	0.235*	-0.217	0.187	0.262*
Top 100 Artists	9.50%	13.86%	0.323	0.101	-0.218	0.074	0.295**
Returns on other assets:							
Stocks	11.25%	17.41%	0.358	-	0.258*	-0.073	0.205
Corporate bonds	8.72%	8.92%	0.415	-	-	-0.278**	-0.144
Commodities	8.63%	23.27%	0.155	-	-	-	0.136
Real estate	4.33%	5.88%	-0.118	-	-	-	-
Treasuries	5.02%	3.09%	-	-	-	-	-

Table 5: Optimal asset allocation

Panel A shows mean-variance tangency portfolio weights on paintings (using either the non-selection-corrected GLS index, the selection-corrected indices A, B, and C from Table 3, or separate selection-corrected indices by style or top 100 artists by sales over the preceding decade), stocks, corporate bonds, commodities, and real estate (as defined in Table 4). The *Benchmark* portfolio excludes art from the investment opportunity set. Panels B and C show the optimal weights for a one-period mean-variance utility investor with risk aversion (γ) equal to two and ten, respectively, where the risk-free asset is the global Treasury bill rates from Dimson, Marsh, and Staunton (2002). Short sales and borrowing are not allowed. Returns are measured over 1961 to 2013. Sharpe ratios are annualized.

	Bench-mark	GLS	Selection-corrected				
			A	B	C	Styles	Styles + Top100
Panel A. Tangency portfolio weights							
Paintings		0.305	0	0	0		
Post-war and Contemporary						0	0
Impressionist and Modern						0	0
Old Masters						0	0
American						0.256	0.217
19th Century European						0	0
Other Styles						0	0
Top 100 Artists							0.048
Stocks	0.288	0.111	0.288	0.288	0.288	0.122	0.118
Corporate bonds	0.522	0.490	0.522	0.522	0.522	0.487	0.486
Commodities	0.190	0.094	0.190	0.190	0.190	0.136	0.131
Real estate	0	0	0	0	0	0	0
Sharpe ratio	0.586	0.644	0.586	0.586	0.586	0.635	0.639

Panel B. Mean-variance utility, risk aversion $\gamma = 2$							
Paintings		0.052	0	0	0		
Post-war and Contemporary						0	0
Impressionist and Modern						0	0
Old Masters						0	0
American						0	0
19th Century European						0	0
Other Styles						0	0
Top 100 Artists							0.092
Stocks	0.607	0.437	0.607	0.607	0.607	0.607	0.559
Corporate bonds	0.234	0.351	0.234	0.234	0.234	0.234	0.215
Commodities	0.159	0.160	0.159	0.159	0.159	0.159	0.134
Real estate	0	0	0	0	0	0	0
Treasuries		0	0	0	0	0	0
Sharpe ratio	0.504	0.562	0.504	0.504	0.504	0.504	0.512
Panel C. Mean-variance utility, risk aversion $\gamma = 10$							
Paintings		0.246	0	0	0		
Post-war and Contemporary						0	0
Impressionist and Modern						0	0
Old Masters						0	0
American						0.155	0.044
19th Century European						0	0
Other Styles						0	0
Top 100 Artists							0.149
Stocks	0.241	0.143	0.241	0.241	0.241	0.183	0.163
Corporate bonds	0.460	0.489	0.460	0.460	0.460	0.485	0.492
Commodities	0.176	0.123	0.176	0.176	0.176	0.167	0.151
Real estate	0	0	0	0	0	0	0
Treasuries	0.124	0	0.124	0.124	0.124	0	0
Sharpe ratio	0.585	0.640	0.585	0.585	0.585	0.614	0.619

Appendix to

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

Arthur Korteweg, Roman Kräussl, and Patrick Verwijmeren*

This appendix presents additional results and robustness checks for the paper "Does it Pay to Invest in Art? A Selection-corrected Returns Perspective." In Section A, we present extended descriptive statistics of buy-ins, and we compare our repeat sales sample to the full BASI database. Section B describes the price index construction. Section C explains the estimation procedure in detail. Section D shows the sale probability function before and after 1998. Section E discusses the selection model estimates when we use buy-in data in addition to the repeat sales. Section F lists the various robustness checks on the portfolio allocation results. Section G shows the importance of the non-linearities in the sale probability function for art returns.

A. Extended descriptive statistics

Table A.1 shows descriptive statistics of the sample of 3,854 buy-ins between 2007 and 2013. Compared to the repeat sales sample, paintings that are bought in tend to be larger and have higher price estimates. Buy-ins happen slightly more often at Christie's and Sotheby's and in London and New York auction houses. Figure A.1 shows the time-series of buy-ins.

Table A.2 shows the descriptive statistics for the full BASI sample over the 1960 to 2013 sample period, compared to the characteristics of the repeat sales sample. Larger and more

* Arthur Korteweg (korteweg@marshall.usc.edu) is from the University of Southern California Marshall School of Business, Roman Kräussl (roman.kraussl@uni.lu) is from the 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, and University of Glasgow.

expensive paintings 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.

B. Price index construction

The art literature defines the price index, $\Pi(t)$, across N artworks relative to the base year 0 as

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

where $P_i(t)$ is the price of painting *i* at time *t*. Goetzmann (1992) and Goetzmann and Peng (2002) show that the index can be constructed from the model estimates in a sequential manner starting from the base year index normalized at 100,

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

where the $\frac{1}{2}\sigma^2$ adjusts for the Jensen inequality term due to the log operator in equation (1) of the main text.

C. Estimation procedure

For each painting, $i = 1 \dots N$, the natural logarithm of price, $p_i(t)$, at time $t = 1 \dots T$, follows the process

$$p_i(t) = p_i(t-1) + \delta(t) + \varepsilon_i(t), \tag{A.3}$$

with $\varepsilon_i(t) \sim \mathcal{N}(0, \sigma^2)$ and i.i.d. across paintings. The price is observed when $w_i(t) \geq 0$, where the latent variable $w_i(t)$ is given by the selection equation,

$$w_i(t) = \sum_{k=1}^{K} g_k(p_i(t); Z_i(t)) \cdot \alpha_k + W_i(t)'\alpha_w + \eta_i(t). \tag{A.4}$$

This equation is linear in its parameters $\alpha = [\alpha_1, \ldots, \alpha_K, \ \alpha_w]'$, but may be non-linear in prices. The vector $Z_i(t)$ contains observed data, for example historical auction prices that can be used to construct returns. The error term $\eta_i(t)$ is distributed i.i.d. $\mathcal{N}(0,1)$ and is uncorrelated with $\varepsilon_i(t)$.

We estimate the set of parameters, $\theta = (\delta(2) \ldots \delta(T), \alpha, \sigma^2)$, using a Bayesian estimation procedure. We augment the parameter set with the latent variables (Tanner and Wong, 1987), and draw a large number of samples from the joint posterior distribution, $f(\theta, \{p_i(t), w_i(t)\}|data)$ using a Gibbs sampler (Gelfand and Smith, 1990). The sampler iteratively draws realizations from the following three complete conditional distributions, leading to a sequence of samples that converges to a sample from the joint posterior distribution:

1. Latent prices: $f(\{p_i(t)\}|\{w_i(t)\}, \theta, data)$.
2. Selection variables: $f(\{w_i(t)\}|\{p_i(t)\}, \theta, data)$.
3. Parameters: $f(\theta|\{p_i(t), w_i(t)\}, data)$.

After dropping the first 10,000 draws to allow the sampler to converge, we compute the marginal posterior distributions of parameters, $f(\theta|data)$, and the price index by numerically integrating over the next 50,000 draws of the joint posterior distribution. The remainder of this section details how to sample from each conditional distribution.

C.1 Latent prices

Sampling the latent prices is a non-linear filtering problem, where (A.3) is the law of motion and (A.4) is the observation equation. We draw $\{p_i(t)\}$ using the single-state updating Metropolis-

Hastings sampler (Jacquier, Polson, and Rossi, 1994). Given the Markov property of (A.3) and suppressing the conditioning on θ and the observed data (for notational simplicity), the posterior distribution of a single state is

$$f(p_i(t)|p_i(t-1), p_i(t+1), \{w_i(t)\}) \propto f\big(w_i(t)|p_i(t)\big) \cdot f\big(p_i(t)|p_i(t-1), p_i(t+1)\big). \tag{A.5}$$

The first component of (A.5) is a normal distribution, as given by the observation equation (A.4). The latter component, by another application of Bayes' law, can be written as

$$f\big(p_i(t)|p_i(t-1), p_i(t+1)\big) \propto f\big(p_i(t)|p_i(t-1)\big) \cdot f\big(p_i(t+1)\big|p_i(t)\big). \tag{A.6}$$

Using the law of motion (A.3) and basic algebra, this distribution simplifies to

$$p_i(t)|p_i(t-1), p_i(t+1) \sim \mathcal{N}\left(\overline{p}, \frac{1}{2}\sigma^2\right), \tag{A.7}$$

where $\overline{p} = \frac{1}{2}(p_i(t-1) + \delta(t)) + \frac{1}{2}(p_i(t+1) - \delta(t+1))$. In other words, conditional on knowing the future and past price realization (and θ), the expected current price from the state evolution is the average of the one-period ahead forecast and the one-period backwards forecast.

Despite the fact that (A.5) is quite easy to evaluate, being the product of two normal probability density functions, it is not a known distribution that can be sampled directly. Therefore, we use a Metropolis-Hastings step (see, for example, Robert and Casella, 2005, for a lucid treatise), in which we first make a proposal for $p_i(t)$ from an easy-to-sample distribution, and then accept or reject it based on how likely this proposal would be generated under the true distribution (A.5). A simple scheme would be to propose from (A.7), but this could lead to poor acceptance properties as no information from (A.4) is incorporated. Instead, we combine (A.7) with a linearized version of (A.4) based on a Taylor expansion around $\overline{p}$,

$$w_i(t) \approx c(\overline{p}) + \sum_{k=1}^{K} \frac{\delta g_k(\overline{p}; Z_i(t))}{\delta p_i(t)} \cdot \alpha_k \cdot (p_i(t) - \overline{p}) + W_i(t)'\alpha_w + \eta_i(t). \tag{A.8}$$

Plugging (A.8) and (A.7) into (A.5) and simplifying terms yields the proposal distribution

$$\hat{f}\big(p_i(t) - \overline{p} | p_i(t-1), p_i(t+1)\big) \sim \mathcal{N}\left(\frac{\frac{1}{2}\sigma^2}{\frac{1}{2}\sigma^2 + \frac{1}{B^2}} \cdot \frac{A}{B}, \frac{\frac{1}{2}\sigma^2 \cdot \frac{1}{B^2}}{\frac{1}{2}\sigma^2 + \frac{1}{B^2}}\right), \tag{A.9}$$

where $A \equiv w_i(t) - c(\overline{p}) - W_i(t)'\alpha_w$, and $B \equiv \sum_{k=1}^{K} \frac{\delta g_k(\overline{p}; Z_i(t))}{\delta p_i(t)} \cdot \alpha_k$. Another way of viewing the proposal distribution is that it is a weighted combination between (A.7) and the approximated distribution of $p_i(t) - \overline{p} \sim \mathcal{N}(A/B, 1/B^2)$ implied by (A.8). As in a standard shrinkage estimator, the weights are proportional to the precision (the inverse of the variance) of each distribution. If (A.8) carries no information about prices (i.e., if $B = 0$), then this distribution reverts to (A.7). The more price information (A.8) contains (i.e., the larger B), the more weight is put on (A.8) relative to (A.7).

Thus, we sample a proposal, $\hat{p}$, for $p_i(t)$ from (A.9) and accept it with probability

$$Q = \min\left(1, \frac{f(\hat{p})/\hat{f}(\hat{p})}{f(p)/\hat{f}(p)}\right), \tag{A.10}$$

where $f(\cdot)$ is the distribution (A.5), and p is the draw for $p_i(t)$ from the most recent iteration of the sampler. If $\hat{p}$ is rejected, then we keep p for the present iteration. Compared to (A.7), the proposal distribution (A.10) is considerably closer to (A.5), resulting in better mixing properties of the estimator. If $\overline{p}$ happens to fall on a discontinuity of $g(\cdot)$, then we revert to using (A.7) for the proposal distribution, though this almost never happens in our specifications.

The acceptance probability (A.10) corrects for the fact that the proposal is an approximation, and does not coincide with (A.5). A nice property of the above sampling scheme is that the approximation is exact in the special case where (A.4) is linear in prices. The proposal distribution (A.9) then coincides with the Kalman filter posterior distribution, and the acceptance probability is always 100%.[1]

C.2 Selection variables

Sampling from the conditional posterior distribution of the selection variables, $f(\{w_i(t)\}|\{p_i(t)\},\theta,data)$, is akin to sampling from the (augmented) posterior distribution of a probit model (Albert and Chib, 1993). Since $\eta_i(t)$ is i.i.d. we may draw each painting-year separately.

When a painting is sold, its price is observed and the posterior distribution of $w_i(t)$ is

$$w_i(t)|p_i(t),\theta,data \sim \mathcal{N}_L\left(\sum_{k=1}^{K} g_k(p_i(t);Z_i(t))\cdot\alpha_k + W_i(t)'\alpha_w, 1\right), \qquad \text{(A.11)}$$

where $\mathcal{N}_L(\mu,\sigma^2)$ is the normal distribution with mean μ and variance σ^2, truncated from below at zero.

Conversely, in periods in which a painting is not sold,

$$w_i(t)|p_i(t),\theta,data \sim \mathcal{N}_U\left(\sum_{k=1}^{K} g_k(p_i(t);Z_i(t))\cdot\alpha_k + W_i(t)'\alpha_w, 1\right), \qquad \text{(A.12)}$$

where $\mathcal{N}_U(\mu,\sigma^2)$ is the normal distribution with mean μ and variance σ^2, truncated from above at zero.

[1] A similar, though slightly more complex proposal distribution, can be found by applying the extended Kalman filter to last period's price to obtain a "forward" distribution of the current price, and applying it backwards to next period's price to get a "backward" distribution. The proposal distribution is a mix of these two normal distributions in proportion to their relative precisions. This proposal also simplifies to the Kalman filter when (A.4) is linear.

C.3 Parameters

Since the function $g(\cdot)$ in (A.4) is linear in α_r, the conditional distributions of α, $\{\delta(t)\}$, σ^2, are determined from the two Bayesian linear regressions in (A.3) and (A.4). These are estimated separately, since the error terms are independent by assumption.

In (A.3), $\delta = [\delta(2), \dots, \delta(T)]'$ and σ^2 are the parameters of a regression of Y_p on X_p. The vector Y_p stacks the one-period log returns, $p_i(t) - p_i(t-1)$, across all paintings and time periods. It is of length $\sum_{t=2}^{T} N(t)$ where $N(t)$ is the number of paintings for which $p_i(t) - p_i(t-1)$ exists. The matrix X_p is a $\sum_{t=2}^{T} N(t)$ by $T-1$ matrix of zeros and ones. Each row has exactly $T-2$ zeros, and a one in column $t-1$, which corresponds to the date of the return in Y_p. For example, the first column of X_p corresponds to $p_i(2) - p_i(1)$.

We use a standard conjugate normal-inverse gamma prior,

$$\sigma^2 \sim IG(a_0, b_0), \tag{A.13}$$

$$\delta|\sigma^2 \sim \mathcal{N}(\mu_0, \sigma^2\Sigma_0^{-1}). \tag{A.14}$$

The posterior distributions are

$$\sigma^2|Y_p, X_p \sim IG(a, b), \tag{A.15}$$

$$\delta|\sigma^2, Y_p, X_p \sim \mathcal{N}(\mu, \sigma^2\Sigma^{-1}), \tag{A.16}$$

where

$$a = a_0 + \sum_{t=2}^{T} N(t), \tag{A.17}$$

$$b = b_0 + e'e + (\mu - \mu_0)'\Sigma_0(\mu - \mu_0), \tag{A.18}$$

$$\Sigma = \Sigma_0 + X_p' X_p, \tag{A.19}$$

$$\mu = \Sigma^{-1}\left(\Sigma_0\mu_0 + X_p' Y_p\right). \tag{A.20}$$

The vector $e = Y_p - X_p\mu$ contains stacked error terms. It is numerically more efficient to directly construct $X_p'X_p$ and $X_p'Y_p$, avoiding large matrix manipulations. This is feasible due to the unique structure of X_p.

The parameters $\alpha = [\alpha_1, \dots, \alpha_K, \ \alpha_w]'$ in (A.4) are estimated from the linear regression of Y_w on X_w with known variance equal to 1 (due to the normalization of $\eta_i(t)$). The vector Y_w stacks the selection variables, $w_i(t)$, across all paintings and time periods. The matrix X_w stacks $[1, g_1(r_i^0), \dots, g_K(r_i^0), W_i(t)']$ over all paintings and time periods.

The prior distribution is

$$\alpha \sim \mathcal{N}(\theta_0, \Omega_0^{-1}), \tag{A.21}$$

and the posterior is

$$\alpha | Y_w, X_w \sim \mathcal{N}(\theta, \Omega^{-1}), \tag{A.22}$$

with

$$\Omega = \Omega_0 + X_w'X_w, \tag{A.23}$$

$$\theta = \Omega^{-1}(\Omega_0\theta_0 + X_w'Y_w). \tag{A.24}$$

C.4 Buy-in information

As discussed in detail in the paper (Sections 1 and 2.4), when a painting is bought in we use the low price estimate as an upper bound on its value. Thus, we treat the price as unobserved in the filtering step in Section C.1, and truncate the proposal distribution (A.9) and the true distribution (A.5) from above at the low price estimate. To draw the latent selection variable in Section C.2 for the buy-in events, we use (A.11) as it reflects the fact that the painting was taken to auction.

C.5 Priors and starting values

We use diffuse priors for all parameters. The prior distribution for σ^2 is inverse gamma with $a_0 = 2.1$, and $b_0 = 10$. Under this prior, $E[\sigma] = 27.0\%$ per year, and σ is between 11.5% and 90.4% with 99% probability.

The prior mean for the change in the log-price index δ is zero (i.e., $\mu_0 = 0$), and we set Σ_0^{-1} equal to the identity matrix. Together with the prior on σ^2, this means that our prior on the annual log-return of the price index is between -97% and +97% with 99% probability.

We set the prior mean for all α parameters equal to zero (i.e., $\theta_0 = 0$), and we set Ω_0^{-1} equal to a diagonal matrix with 10,000 on the diagonal and zeros elsewhere, implying that each component of α is between -258 and +258 with 99% probability.

We start the algorithm with α and δ equal to zero, and $\sigma^2 = 16\%$. We do not need starting values for the unobserved prices, as they are filtered out in the first step, before they are needed elsewhere. We do not need starting values for the selection variables because with $\alpha = 0$, the unobserved prices do not depend on w in the very first cycle of the algorithm.

D. Split-sample estimates

There has been a shift in the accessibility of auction records, which gradually became more widely available in the mid to late 1990s. With potential buyers and sellers having more and better data in the last 15 years of the sample, this could have an impact on the sale probability function. It is difficult to pin down an exact date when auction records became more easily accessible. We chose to split the sample into the periods before and after 1998, which is the year when Christie’s and Sotheby’s started providing online provenance information on all auction sales.

Table A.3 shows the coefficients of the sale probability functions in both sub-samples. The discontinuity and V-shape are present in both sub-samples. The two most notable differences are that pre-1998, the slope on gains is higher, and the non-linear component on losses is weaker. How much of these differences are due to changes in database access, versus changes in the market or the macro-economy is difficult to say, though. Still, it is reassuring that the general shape of the sale probability function is the same in both sub-samples.

E. Estimates from sample of repeat sales and buy-ins

Table A.4 shows the selection model parameter estimates for the sample of repeat sales augmented with buy-in data. Compared to Table 3 in the paper, the selection is slightly stronger when paintings have increased in value since their prior sale. This pushes the index lower, as can be seen in Panel B of Figure 3 in the paper.

The reason for this result is somewhat subtle, as the inclusion of buy-ins changes the dynamics of a painting's price process. To the econometrician, the price information in a buy-in is both positive and negative. On one hand, the fact that a buy-in occurred reveals that the owner had a positive price signal. On the other hand, the fact that the painting did not get sold is negative information, as the price must be lower than the lowest price estimate. The latter channel dominates, and thus we see lower index values when buy-ins are considered.

F. Robustness of portfolio allocation results

In this section we confirm that the exclusion of a broadly diversified portfolio of art from optimal portfolios after correcting for sample selection, is robust to various realistic extensions.

First, when using the index returns that include buy-ins for the 2007 to 2013 part of the sample, we find portfolio allocations that are very close to the results in the paper. This is not surprising as we only have six years of buy-in information. With a longer time series the Sharpe ratios on art may change more dramatically. Given that art returns are likely even lower when including buy-ins, this would reaffirm the unattractiveness of a broadly diversified art portfolio.

Second, Ang, Papanikolaou, and Westerfield (2014) show that allocation to an illiquid asset is lower when the illiquid nature of the asset is taken into account. For example, using their model[2] and our non-selection-corrected returns, a power utility investor with risk aversion of ten allocates 9% of her portfolio to art, down from the 25% allocation in Table 5 of the paper. Not surprisingly, the allocation to a broad art portfolio after correcting for selection remains at zero for all levels of risk aversion.

Third, instead of the Dimson (1979) method that we use in the paper to correct for spurious autocorrelation in returns, some papers in the art-investments literature (e.g., Campbell, 2008, and Renneboog and Spaenjers, 2013) use the technique pioneered in real estate by Geltner (1991) to unsmooth the index. Using this alternative method does not change the conclusions.

Fourth, if investors have, for example, power utility rather than mean-variance utility, they may care about higher moments of returns such as skewness (e.g., Kraus and Litzenberger, 1983, Ball, Kothari, and Shanken, 1995, and Harvey and Siddique, 2000) and kurtosis (e.g., Dittmar, 2002). Art returns are slightly positively skewed (the skewness of the art index returns is 0.72, versus -0.74 for stocks), which is attractive. However, the kurtosis of art returns is 4.52, which is considerably higher than the 3.27 kurtosis of stocks. On balance, the higher moments do not

---

[2] We thank Andrew Ang, Dimitris Papanikolaou, and Mark Westerfield for generously sharing their code for the portfolio optimization problem in their paper.

make art particularly more attractive relative to stocks, as is borne out by the Ang et al. model results, which assumes power utility.

Fifth, there are sizeable transaction costs in the art market. The typical buyer's premium in art is up to 17.5 percent of the hammer price, and on top of that there are storage and insurance fees. These costs make paintings less appealing for optimal portfolio allocation and hence reinforce our main result on broadly diversified portfolios.

Finally, our portfolio allocation results are also robust to using logarithmic rather than arithmetic returns, and to dropping short-duration repeat sales that occur within one year or less after the previous sale.

G. Impact of non-linearities in the sale probability function on art returns

We compare the art returns from the non-linear sale probability Model A from the paper to a linear sale probability model as in Korteweg and Sorensen (2010, 2014), estimated on the same sample of paintings (note that the returns from the other non-linear models in the paper are very close to Model A, and for brevity will not be reported here). The average annual art return from the linear model is 6.67 percent, compared to 6.29 percent for Model A. The annual standard deviation is 12.93% in the linear model, whereas it is 11.42% in the non-linear model. This results in an annual Sharpe ratio of 0.128 versus 0.111 for the linear and non-linear model, respectively. By the end of the sample period, the linear model overstates the price index by 11.2%, relative to the non-linear model. These results underscore the importance of accounting for non-linearities when estimating art indices and returns.

References

Albert, J., and Chib, S., 1993. Bayesian analysis of binary and polychotomous response data. *Journal of the American Statistical Association* 88, 669-679.

Ang, A., Papanikolaou, D., and Westerfield, M., 2014. Portfolio choice with illiquid assets. *Management Science* 60, 2737-2761.

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.

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.

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

Dittmar, R., 2002. Nonlinear pricing kernels, kurtosis preference, and evidence from the cross-section of equity returns. *Journal of Finance* 57, 369-403.

Gelfand, A., and Smith, A.F.M., 1990. Sampling-based approaches to calculating marginal densities. *Journal of the American Statistical Association* 85, 398-409.

Geltner, D.M., 1991. Smoothing in appraisal-based returns. *Journal of Real Estate Finance and Economics* 4, 327-345.

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

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

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

Jacquier, E., Polson, N.G.P., and Rossi, P.E., 1994. Bayesian analysis of stochastic volatility models. *Journal of Business and Economic Statistics* 94, 371–417.

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., 2014. Estimating loan-to-value distributions. *Working paper*, Stanford University and Columbia University.

Kraus, A., and Litzenberger, R., 1983. On the distributional conditions for a consumption-oriented three moment CAPM. *Journal of Finance* 38, 1381-1391.

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

Robert, C., and Casella, G., 2005. Monte Carlo Statistical Methods. 2nd ed. New York: Springer.

Tanner, M., and Wong, W., 1987. The calculation of posterior distributions by data augmentation. *Journal of the American Statistical Association* 82. 528-549.

Figure A.1: Number of buy-in sales

This figure shows the time-series of the number of buy-ins in the sample, where paintings went to auction but did not sell.

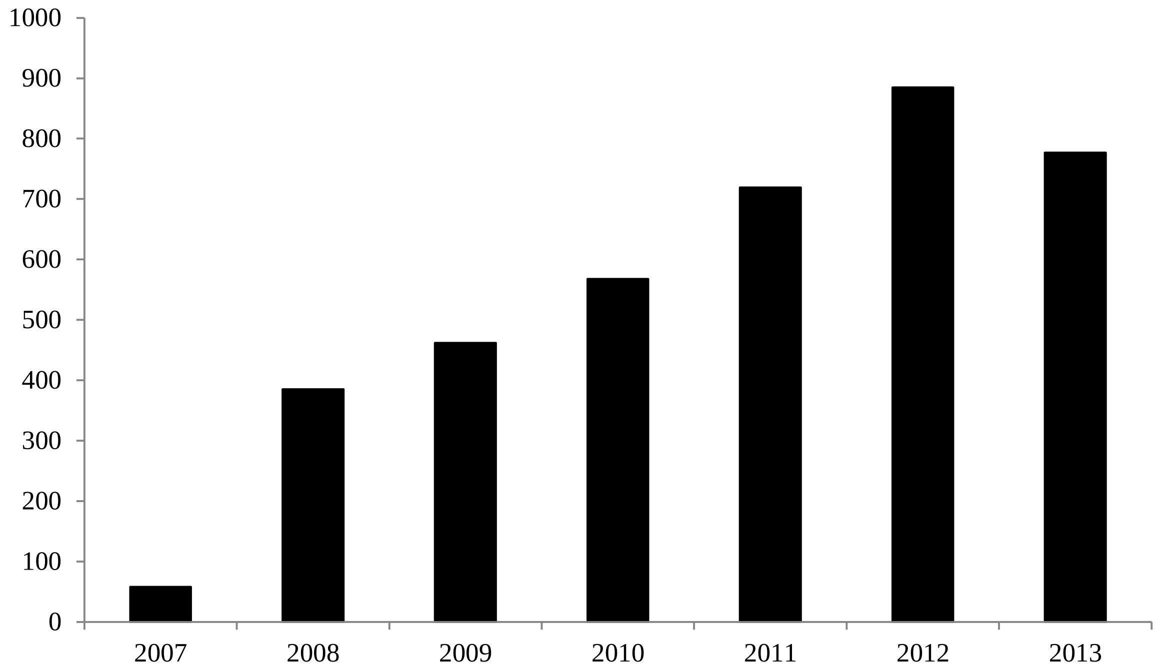

Table A.1: Descriptive statistics for buy-ins

This table reports descriptive statistics for the sample of 3,854 auctions that did not result in a sale ("buy-ins") in the Blouin Art Sales Index (BASI) dataset from 2007 to 2013. *Low Estimate* and *High Estimate* are the auction house's low and high price estimates, respectively, in thousands of U.S. dollars. *Surface* is the surface of the painting in thousands of squared millimeters. *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 the decade prior to the year of sale, and zero otherwise. The remaining variables represent style classifications.

	Mean	Median	St. Dev.
Low estimate ($000s)	220.0	39.8	1,081.0
High estimate ($000s)	310.4	55.3	1,477.5
Surface	810.4	437.5	1,365.1
Deceased < 2yrs	1.27%		
Christie's	42.63%		
Sotheby's	41.10%		
London	34.20%		
New York	35.68%		
Post-war and Contemporary	14.95%		
Impressionist and Modern	20.81%		
Old Masters	17.51%		
American	7.47%		
19^{th} Century European	20.58%		
Other Style	18.68%		
Top 100 Artists	12.30%		

Table A.2: Full BASI sample comparison

This table compares descriptive statistics for the full sample and the repeat sales sample of paintings in the Blouin Art Sales Index (BASI) dataset from 1960 to 2013. The table 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. *Low Estimate* and *High Estimate* are the auction house's low and high price estimates, respectively, as a percentage of the hammer price. *Surface* is the surface of the painting in thousands of squared millimeters. *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 the decade prior to the year of sale, and zero otherwise. The remaining variables represent style classifications. The last column shows *t*-statistics for difference in means tests (for continuous variables) and *z*-statistics for difference in proportions tests (for dummy variables) between the full and the repeat sales samples. ***, ** and * indicate statistical significance at the 1, 5 and 10 percent level, respectively.

	Repeat sales sample (69,103 sales)			Full sample (2,715,300 sales)			Difference statistic
	Mean	Median	St. Dev.	Mean	Median	St. Dev.	
Hammer price ($000s)	150.6	14.5	923.8	31.5	2.8	457.6	65.14***
Low estimate (% of hammer price)	85.64%	81.63%	76.50%	88.02%	83.33%	63.06%	-6.72***
High estimate (% of hammer price)	118.83%	114.29%	83.23%	125.35%	115.38%	107.66%	-9.98***
Surface	630.3	352.3	1,161.1	500.5	306.5	805.4	40.81***
Deceased < 2yrs	1.72%			1.74%			-0.41
Christie's	32.58%			14.94%			126.98***
Sotheby's	32.67%			15.30%			123.80***
London	26.87%			13.76%			97.80***
New York	29.11%			9.18%			175.11***
Post-war and Contemporary	14.32%			11.11%			26.38***
Impressionist and Modern	25.80%			14.31%			84.54***
Old Masters	11.23%			9.73%			13.17***
American	10.23%			7.64%			25.15***
19^{th} Century European	21.20%			30.46%			-52.33***
Other Styles	17.22%			26.75%			-56.07***
Top 100 Artists	15.20%			3.15%			171.30***

Table A.3: Sample splits

This table presents the parameter estimates of the selection equation (Equation (3) in the paper), estimated on the sub-sample of auction data before and after 1998. *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. *Sigma* is the standard deviation of the error term in Equation (1) of the paper. Standard errors are in parentheses. The column *Difference p-value* contains the p-value for the test with null hypothesis that the coefficients are the same in the two sub-samples. ***, ** and * indicate statistical significance at the 1, 5 and 10% level, respectively.

	Sample		Difference
	<1998	>=1998	p-value
Return > 0	0.119***	0.113***	0.825
	(0.022)	(0.016)	
(Return<0)	-0.610***	-0.652***	0.675
* return	(0.081)	(0.062)	
(Return<0)	-0.038	-0.462***	0.000
* return^2	(0.074)	(0.048)	
(Return>0)	0.280***	0.072**	0.000
* return	(0.040)	(0.031)	
(Return>0)	0.229***	0.258***	0.356
* return^2	(0.023)	(0.022)	
Time (years)	-0.030***	-0.007***	0.000
	(0.003)	(0.001)	
Time squared	0.000**	0.000**	0.076
	(0.000)	(0.000)	
Style fixed effects	Y	Y	
Sigma	0.200***	0.200***	
	(0.001)	(0.001)	

Table A.4: Selection equation coefficients, with buy-ins

This table presents the parameter estimates of three specifications of the selection equation (Equation (3) in the paper), using the sample of repeat sales supplemented with buy-ins. R*eturn* is the natural logarithm of the return since the prior sale of a painting. *Relative share* is the market share (in terms of sales) of the painting's style in the year of sale compared to the style's average market share in the five years prior to the sale. *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 the Historical Statistics of the World Economy. The other variables are as defined in Table 1 of the paper. *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% level, respectively.

	A		B		C	
Return > 0	0.124	***	0.122	***	0.047	
	(0.013)		(0.013)		(0.049)	
(Return > 0)					0.075	*
* relative share					(0.047)	
(Return<0) * return	-0.644	***	-0.652	***	-0.601	***
	(0.048)		(0.048)		(0.150)	
(Return<0) * return					-0.049	
* relative share					(0.139)	
(Return<0) * return^2	-0.393	***	-0.395	***	-0.412	**
	(0.037)		(0.037)		(0.156)	
(Return<0) * return^2					0.018	
* relative share					(0.148)	
(Return>0) * return	0.156	***	0.159	***	0.507	***
	(0.024)		(0.024)		(0.139)	
(Return>0) * return					-0.347	**
* relative share					(0.137)	
(Return>0) * return^2	0.232	***	0.233	***	0.067	
	(0.015)		(0.015)		(0.079)	
(Return>0) * return^2					0.167	**
* relative share					(0.077)	
Relative share					-0.010	
					(0.020)	
Time (years)	-0.014	***	-0.014	***	-0.014	***
	(0.001)		(0.001)		(0.001)	
Time squared	0.000	***	0.000	***	0.000	***
	(0.000)		(0.000)		(0.000)	
Log (surface)			0.010	***	0.010	***
			(0.003)		(0.003)	
Deceased < 2 years			0.039		0.040	
			(0.027)		(0.027)	

World GDP growth			1.862	***	1.859	***
			(0.194)		(0.193)	
Style fixed effects	Yes		Yes		Yes	
Sigma	0.201	***	0.201	***	0.201	***
	(0.001)		(0.001)		(0.001)	