Management Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Speculative Trading and Bubbles: Evidence from the Art Market

Julien Pénasse, Luc Renneboog

To cite this article:
Julien Pénasse, Luc Renneboog (2022) Speculative Trading and Bubbles: Evidence from the Art Market. Management Science 68(7):4939-4963. https://doi.org/10.1287/mnsc.2021.4088

informs
http://pubsonline.informs.org/journal/mnsc

MANAGEMENT SCIENCE
Vol. 68, No. 7, July 2022, pp. 4939–4963
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

Speculative Trading and Bubbles: Evidence from the Art Market

Julien Pénasse,[a] Luc Renneboog[b]

[a] University of Luxembourg, L-1359 Luxembourg, Luxembourg; [b] Tilburg University, 5000 LE Tilburg, Netherlands

Contact: julien.penasse@uni.lu, https://orcid.org/0000-0003-3738-4549 (JP); luc.renneboog@uvt.nl, https://orcid.org/0000-0002-6068-3910 (LR)

Received: March 6, 2018
Revised: December 1, 2020; February 21, 2021
Accepted: March 17, 2021
Published Online in Articles in Advance: September 7, 2021

https://doi.org/10.1287/mnsc.2021.4088

Abstract. We argue that extrapolative expectations drive boom–bust cycles in the postwar art market. Price run-ups coincide with increases in demand fundamentals but are followed by predictable busts. Predictable changes account for about half of the variance of five-year price changes. High prices coincide with many attributes of speculative bubbles: trading volume, the share of short-term trades, the share of postwar art, and volatility are all higher during booms. In addition, short-term transactions underperform long-term transactions. Survey evidence further confirms the link between beliefs, prices, and volume dynamics as in models in which extrapolative beliefs fuel speculative bubbles.

History: Accepted by Tyler Shumway, finance.

Supplemental Material: The online appendix is available at https://doi.org/10.1287/mnsc.2021.4088.

Keywords: extrapolative beliefs • speculative bubbles • trading volume • art auction

"As a collector, I trade all the time, it's the capitalist in me." (Peers 2008)[1]

1. Introduction

The history of financial markets is replete with episodes of booms and busts in asset prices that are difficult to reconcile with underlying economic fundamentals. To explain these episodes, theory offers a number of mechanisms, including excessive credit, over-leverage, and the destabilizing role of speculation (see, e.g., Brunnermeier and Oehmke 2012, Xiong 2013). We present evidence that biased expectations play an important role in driving boom–bust cycles in the postwar art market. Art price booms coincide with periods of strong demand for art but tend to be followed by predictable busts. This property is central to models in which investors over-extrapolate past demand: investors overbid when demand is booming (and underbid when demand is falling) because they over-extrapolate demand trends into the future.

Figure 1 illustrates our main findings. Panel (a) plots aggregate prices, and panel (b) compares five-year growth in top 1% U.S. wealth with the five-year price growth, measured over the subsequent five years. We interpret top wealth as an indicator of art demand (Goetzmann et al. 2011) and price changes as the return to investing in a diversified art portfolio. We see that art market booms and busts coincide with parallel cycles in demand.

Because the supply of art is inelastic, it is tempting to conclude that demand cycles cause art price booms and busts. However, variation in demand does not imply anything about future returns. We see in panel (b) that, on average, demand booms are followed by low, often negative returns. For instance, the year 1990—a year characterized by buoyant prices and surging wealth—was followed by the most dramatic bust in the postwar era, an episode that is often described as the late 1980s bubble. Five years earlier, although prices were stagnant and wealth growth had reached an all-time low, investors could have expected to make large profits by investing in art. The same inverse relationship repeatedly reappears in our sample (the correlation between the two series in panel (b) is a whopping 0.54).

This lead–lag relation between wealth and returns is at odds with textbook theory in which asset prices are forward-looking. In rational expectation models with constant risk premiums, prices fluctuate to reflect investors' expectations about future dividends. In the context of art, dividends correspond to the rent collectors are willing to pay to own art at a given point in time. As supply is inelastic, the rent is determined by

Figure 1. (Color online) Past Wealth Predicts Future Price Changes

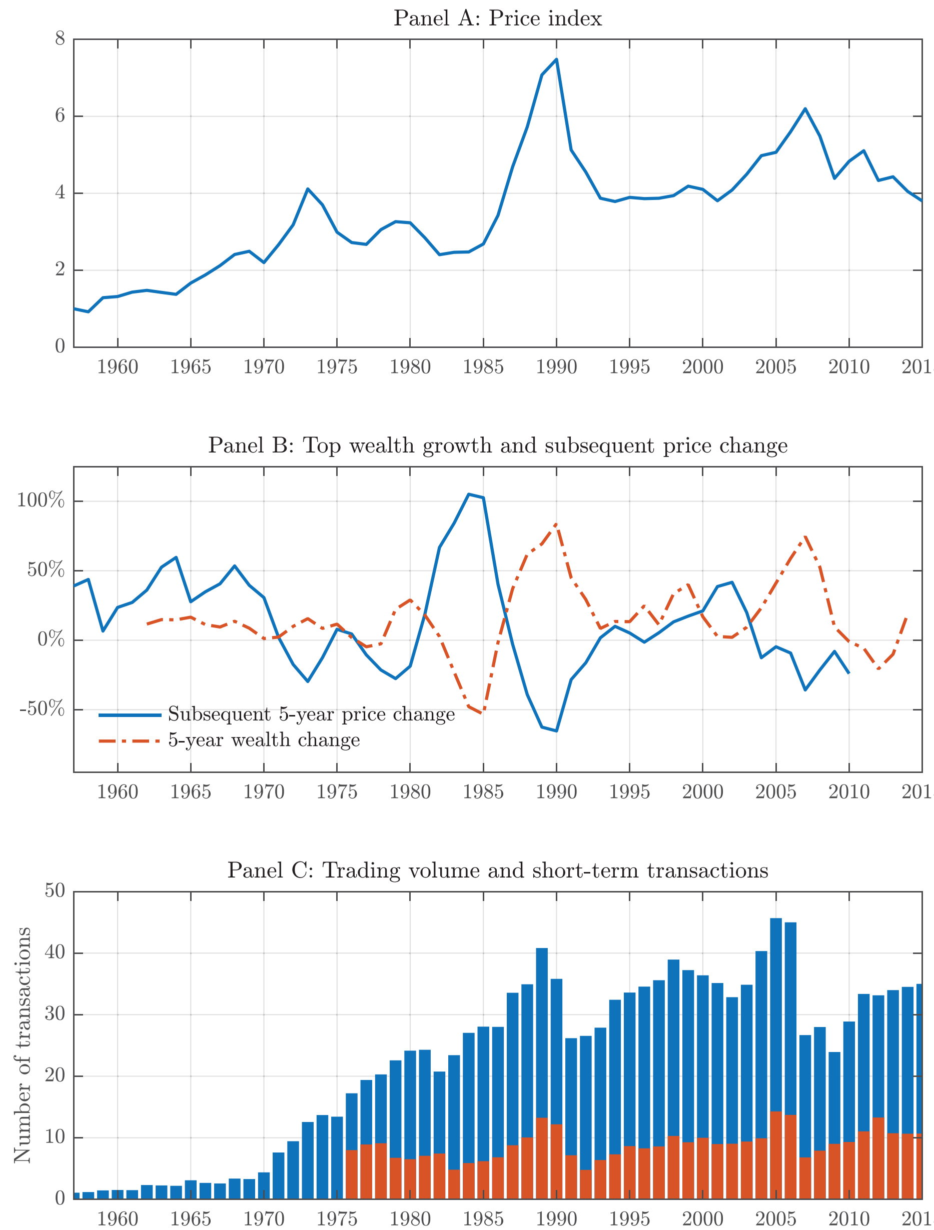

Notes. Panel (a) shows the aggregate hedonic price index. Panel (b) shows the five-year wealth growth rate (measured as the change in a five-year trailing average of the wealth held by the 1% top earners) and subsequent aggregate change in art prices (also measured over five years). Panel (c) shows the total number of transactions (thousands) in blue (dark grey) as well as an estimate of the number of transactions with holding periods less than two years (red/light grey).

the demand for art. Textbook theory, thus, predicts that art returns lead wealth growth, whereas we find the opposite.

We argue that wealth growth predicts future returns because investors have biased expectations. Our interpretation works as follows. Wealth growth is negatively autocorrelated in the data, a feature that surprises investors. When wealth is booming, investors expect wealth will continue growing into the future, bidding prices up. Price booms are followed by predictable busts that occur once investors realize that the demand boom they expected to persist does not materialize. The idea that market participants extrapolate exogenously given cash flows is related to the well-known representativeness heuristic in which subjects draw strong conclusions from small samples of data (Kahneman and Tversky 1974). Extrapolative expectations are increasingly common in behavioral finance models (e.g., Greenwood and Hanson 2015, Barberis et al. 2018, Defusco et al. 2020 for recent examples). We find little evidence for alternative explanations in which investors' expectations are unbiased.

For instance, a time-varying risk premium explanation would require the risk premium to vary dramatically over time in a way that is disconnected from popular risk premium proxies.

If some investors make predictable mistakes, why don't rational investors step in and push back prices to fundamentals? One explanation is that it is essentially impossible to sell short art. Consequently, when prices exceed demand fundamentals, rational investors cannot do more than sell their art to optimist investors. Figure 1 (panel (c)) shows that trading volume—defined as the total number of transactions—is larger during booms, which could reflect sell-offs by rational investors.

In fact, many historical accounts of financial bubbles, such as the South Sea Bubble of 1720 or the stock market boom of the late 1920s, found that high valuation levels often coincided with heightened trading volume (Galbraith 1954, Carlos et al. 2006). The term "speculative bubble" conjures the idea of a market timing game in which investors buy overvalued assets hoping to sell to a greater fool before the crash. Several theories link differences of opinion with trading volume, in particular, during bubbles. The models in Scheinkman and Xiong (2003) and the subsequent literature formalize the idea that, when combined with short-sale constraints, differences of opinion can lead to significant price bubbles together with trading frenzies. In this class of models, investors trade aggressively, not only because they disagree, but also because they understand that they will disagree again in the future. In fact, even rational investors may decide it is more profitable for them to ride a bubble (Abreu and Brunnermeier 2003).

To better understand trading volume, we study its underlying components. Examples abound of art speculators buying art with the intention of swiftly reselling it for a quick profit. Panel (c) indicates that this practice of "flipping" art is more common during booms. We observe a more than proportional increase of short-term transactions (defined as transactions with holding periods less than two years) during the late 1980s bubble, a property that holds over the whole sample. Interestingly, Bayer et al. (2020) and Defusco et al. (2020) document similar patterns in the housing boom of the 2000s, which they interpret as evidence of speculative trading. Anecdotal evidence also suggests that speculators tend to concentrate on the most "glamorous" styles. For instance, we find that, during booms, volume in modernist postwar and contemporary art indeed rises much more than for other styles.

Furthermore, despite higher market volumes, art prices tend to be more volatile during booms. This fragility is another characteristic of many historical bubbles such as the South Sea Bubble, and more recently, the internet bubble in the late 1990s (see, e.g., Cochrane 2003, Ofek and Richardson 2003).

Our evidence suggests extrapolative beliefs and speculative trading are two key elements of speculative dynamics in the art market. We next investigate if these two elements are connected. One recent model, in particular, proposes that disagreement arises endogenously because some, but not all, investors hold extrapolative expectations (Barberis et al. 2018). Using survey data about household and art investors' expectations, we construct measures of returns expectations and belief dispersion, which enables us to test for a link between investors' expectations, trading, and price dynamics. We find that changes in disagreement are positively correlated with changes in volume and price dispersion as predicted by models of speculative trading, such as Scheinkman and Xiong (2003). In addition, average beliefs and belief dispersion are positively related to lagged returns as well as lagged fundamentals, consistent with the evidence in Piazzesi and Schneider (2009) and the model in Barberis et al. (2018).

The extrapolative models of Barberis et al. (2018) and Defusco et al. (2020) offer additional predictions about price and volume dynamics during bubbles. In Barberis et al. (2018), some of the investors extrapolate past returns. Investors disagree about how to interpret past returns and, thus, trade against one another. This translates into a positive correlation between volume and past returns during bubbles, which disappears after bubbles burst. Defusco et al. (2020) assume that investors have homogenous beliefs but differ in their expected investment horizons. Their model predicts that volume peaks before prices during bubbles because optimists increase their reserve prices in spite of falling trading volume. Although these models speak of the stock and real estate markets, the late 1980s bubble provides an out-of-sample setting to test these model predictions. Strikingly, the data supports both models' predictions, which suggests that both mechanisms are present in the data.

In the final empirical section of the paper, we document new facts about short-term transactions in hope of learning about the motives of short-term buyers. Economists have long debated whether speculators play a correcting force in asset markets. In the models of speculative trading that we sketched, speculators play a destabilizing role, but many models assume the contrary. Lovo and Spaenjers (2018), in particular, propose that art speculators are a stabilizing force in bad times. We, thus, investigate if short-term transactions reflect the trade of informed or, at least, sophisticated investors who could play a stabilizing role. We find that holding periods are shorter for cheaper artworks produced by lesser-known artists and are sold in less prominent auction houses. Holding periods are also shorter for the work of living and postwar artists (i.e., glamour art), which is in agreement with the evidence that glamour stocks tend to exhibit higher volume than value stocks (Hong and Stein 2007).

We next study the relative performance of short-term transactions. We find that short-term trades outperform, on average, before transaction costs in terms of both market timing and selecting artworks that appreciate faster in the cross-section. To mitigate concerns of reverse causality, we show that, on average, a buyer earns lower returns if the item was previously held over a short period of time (e.g., it was flipped). However, the apparent outperformance reverses when taking into account realistic transaction costs in contrast to the model prediction of Lovo and Spaenjers (2018). Furthermore, the returns of short-term transactions are much more volatile than those of long-term transactions. We conclude that short-term traders are not necessarily unsophisticated, but that they perform poorly after transaction costs, which is in line with prior evidence in the stock market (Barber and Odean 2000).

This paper relates to the literature documenting historical evidence on asset price formation and bubbles and to theoretical models on (over)pricing, volume, and disagreement. Previous research provides empirical evidence linking overpricing and volume in dual-class shares markets (Mei et al. 2009) as well as during the Chinese warrant bubble (Xiong and Yu 2011). We establish overpricing in the art market through predictive regressions, echoing Greenwood and Hanson's (2015) evidence in the dry bulk shipping industry. We contribute by providing evidence connecting extrapolative expectations, speculative trading, and bubbles in the art market. The art market exhibits characteristics that differ from other markets: it features unlevered and wealthy investors, financial and technological innovations are largely absent, and transaction costs are substantial. These distinct features make it an interesting laboratory to test theories of asset prices and financial bubbles, which were designed with other markets in mind. In this vein, we also relate to the burgeoning literature using transaction data to study the role of short-term traders in the IPO and housing markets (Krigman et al. 1999, Ellul and Pagano 2006, Chinco and Mayer 2016, Defusco et al. 2020, Bayer et al. 2020). Historical accounts of financial bubbles often emphasize the importance of uncertainty (Kindleberger 1978, Greenwood and Nagel 2009, Pastor and Veronesi 2009) as well as the availability of cheap credit, leading to increased risk taking and high leverage (e.g., Brunnermeier 2009, Favilukis et al. 2016, Di Maggio 2017). Such mechanisms are unlikely to play a role in art markets because investors rarely borrow to buy art. Empirically, we find that credit booms tend to be followed by low returns, which suggests prices are excessive during booms, in line with our extrapolative belief explanation. Taken as a whole, the evidence that we present depicts a coherent picture in which extrapolative investors are too optimistic in booms and too pessimistic in busts, fueling speculative bubbles, in line with the models of Scheinkman and Xiong (2003), Defusco et al. (2020), and Barberis et al. (2018).

This paper is also related to a growing literature that examines the art market employing finance tools. Several studies provide evidence that past art returns can help predict future returns (Cutler et al. 1991, Pesando 1993, Goetzmann 1995, Kräussl et al. 2016). We find that wealth growth subsumes past returns in predicting art returns and contributes to most art return variance. Previous studies also document that art investors sometimes have biased expectations. Mei and Moses (2005) show that high estimates at the time of purchase are associated with adverse subsequent abnormal returns, which suggests that credulous collectors are likely to be influenced by biased presale estimates. Beggs and Graddy (2009) and Graddy et al. (2014) provide evidence of anchoring and loss aversion in art auctions. We show that extrapolative expectations can explain aggregate price fluctuations. Three papers relate art prices to demand fundamentals. Mandel (2009) shows that a conspicuous consumption motive can quantitatively rationalize the demand for art in spite of low ex ante returns. Goetzmann et al. (2011) document long-run relations between art prices and top incomes and wealth. Lovo and Spaenjers (2018) propose a dynamic auction model in which rational agents trade art. Their model predicts that prices and volume increase in good times when demand fundamentals are strong. However, because agents are strictly rational, there are no bubbles, and in particular, strong fundamentals do not predict the negative returns we document empirically.

The remainder of this paper is structured as follows. In the next section, we discuss institutional details and present our data. We argue that top wealth is a good proxy for demand fundamentals, complementing the evidence in Goetzmann et al. (2011). In Section 3, we show that art price booms and busts are predictable. In Section 4, we document new stylized facts about volume and volatility during booms and busts, which we relate to investor beliefs. Section 5 studies short-term transactions in the cross-section. Section 6 maps our evidence to existing theories. Section 7 concludes. An online appendix contains a data appendix, additional results, and robustness checks.

2. Background and Data

2.1. The Art Market

Works of art are usually sold through two types of intermediaries: dealers (galleries) and auction houses.

Dealers mostly serve the primary market and place the work of living artists. Auction houses effectively act as brokers and operate most of the secondary market for art and other luxury goods. According to the Tefaf Art Market Report 2017, the art market is worth about €45 billion worldwide, typically split equally between dealers and auction houses (Pownall 2017). Dealers are small businesses by nature, lightly regulated, so that this market is largely opaque. As a result, dealer prices tend not to be reliable or easily obtainable. We, therefore, work with auction data, which is both reliable and publicly available and has been used to study a broad number of finance and economics questions (e.g., Galenson and Weinberg 2000, Mei and Moses 2005, Beggs and Graddy 2009).

The way the auction market operates has changed little over the centuries. Sellers consign works of art to an auction house, and works are grouped to create an auction. Every auction comes with a catalog listing all of the lots in the sale with information about each piece and features a public presale exhibition. At the auction itself, each item is bid on, one at a time. When an item does not attract any bids or never reaches the reserve price, it is bought in. If it is sold, the winner is required to pay the hammer price along with the buyer's commission and any taxes. The auction house also takes a seller's commission, which is deducted from the hammer price, and passes on the remainder of the proceeds to the consignor. Recently, Sotheby's and Christie's have begun offering loans against art with clients putting up paintings or sculptures as collateral, but this practice is uncommon and confined to major purchases (Thompson 2009). Another recent trend is auction houses providing guarantees to sellers who are concerned that not enough bidders will enter the auctions for their items.[2]

Transaction costs are substantial in the art market. Auction houses typically charge commissions of around 10% for both parties (buyers and sellers), and they may also charge mandatory insurance and shipping expenses (see Pesando 1993, Ashenfelter and Graddy 2003). If a specific lot does not sell, a consignor often faces "buy-back" fees to reimburse the auction house for its various costs, such as photographing, researching, and cataloging the item. Sellers also incur some charges, such as shipping and handling costs, even when a work of art goes unsold (i.e., it is bought in). Art buyers also have to take into account the various costs incurred to maintain, store, restore, and insure works of art, which are nonnegligible.

These large transaction costs suggest that collectors put an object up for sale because they are forced to. As the adage says, art selling is induced by the three Ds: divorce, debt, and death. Yet art-as-investment is increasingly popular. The first formal art fund was launched by the British Rail Pension Fund in 1974, investing in some $70 million worth of art and luxury furniture (Trucco 1989). Fifty or so similar investment funds were launched in the 1970s and 1980s. Among the few funds that survived, performance has been tepid, much weaker than traditional asset classes. The Art & Finance Report 2014, a joint report by Deloitte Luxembourg and ArtTactic, identifies 72 art funds in operation in 2014, among which 55 invest only in Chinese Art, arguably one of the most speculative segments of the market. Art funds are, of course, only the tip of the art investment iceberg. Seventy-six percent of art buyers still view their acquisitions as investments, intending to at least avoid negative returns (Picinati di Torcello and Petterson 2014). In today's world, the number of art advisors amounts to about 300,000 (Reyburn 2014).

2.2. Auction Data and Price Indices

Our primary data sources are the Hislop Art Sales Data and the Blouin Art Sales Index, online databases commonly used in research on art markets. Our database spans the period 1957–2015 and comprises 1,329,330 auctioned paintings (more information about sample selection and the data set is provided in Online Appendix A.1). The mean sales price over all observations for 2015 is US$182,784, and the median transaction price for the same year equals US$9,442.

Two approaches are popular to estimate the returns on infrequently traded goods, such as houses or collectibles. The hedonic regression approach relates the natural logs of real U.S. dollar hammer prices to year dummies while controlling for a wide range of "hedonic" characteristics:

$$p_{i,t,k} = p_t + \alpha_k + X_i\gamma + \varepsilon_{i,t}, \tag{1}$$

where $p_{i,t,k}$ denotes the auction price of artwork i for artist k in year t, p_t and α_k are year and artist fixed effects, and X_i is a vector of observable characteristics that affect the price of art, such as information regarding the work itself (through the inclusion of variables capturing attribution, authenticity, medium, size, and topic) and the sale (through the inclusion of auction house dummies and the auction season (month of sale)). The quantity of interest here is p_t, the hedonic price index, which captures the average price of artworks sold in a given year, controlling for any other observable characteristics (Renneboog and Spaenjers 2013).

The other popular approach available to study returns on infrequently traded assets is working with items that sell more than once. When a work of art is sold twice, the return on investment is directly measurable. This is useful to measure art returns on a more granular level although the most common use of repeat sales is to construct price indices. Formally, taking a first difference of Equation (1) gives

$$\Delta p_{i,t\to t+h} = \Delta p_{t\to t+h} + \Delta\varepsilon_{i,t\to t+h}, \tag{2}$$

where the log-difference is taken between the year of purchase t and the year of sale $t+h$. Hence, $\Delta p_{i,t\to t+h}$ is the log-return for holding a work of art i over the holding period h, and $\Delta p_{t\to t+h}$ denotes log changes in aggregate prices over the same period, that is, the repeat-sale index (see Case and Shiller 1987). In our sample of repeated sales, the average round-trip real return is 13.4% (volatility 81.2%) over an average holding period of 7.2 years.[3]

Although the advantages of repeat-sale transactions are clear in terms of direct return calculation, our repeat-sale sample may not be representative of art returns in general because it is based on a subsample of transactions that reappeared at least once at auctions. We note in Online Appendix A.1 that the hedonic and repeat-sales price indices exhibit very similar dynamics, which suggests selection bias is unlikely to matter for the key results in this paper. Nevertheless, whenever possible, we work both with the full sample of transactions—that is not subject to selection bias—and with the repeat-sale sample.

2.3. Art Prices and Fundamentals

In the absence of a large rental market, the fundamental value of any work of art is difficult to pin down and largely relies on past prices and comparable assets. In this context, auction houses play a crucial role in producing information, usually by preparing and disseminating catalogs and by providing estimates of the works of art that will be sold. Yet evidence indicates that auctioneers inflate these estimates to increase transaction prices (Mei and Moses 2005). Prior research also finds evidence of anchoring in art prices (Beggs and Graddy 2009, Hong et al. 2015) as well as violations of the law of one price (Pesando 1993). Narratives of booms and busts in the art market sometimes mention "manias" that concentrate on specific segments of the market. For instance, Hiraki et al. (2009) document the influx of Japanese collectors who massively bought Impressionist and Post-Impressionist art after the real estate booms in the late 1980s.

Although this body of research may present the art market as idiosyncratic and segmented, we observe substantial comovement across very different segments of the market. For instance, the average correlation across style indices is 0.70. We present these results and document a similar pattern for trading volume in Online Appendix B. The evidence suggests that the bulk of prices and volume movements across various market segments have a common source.

A natural candidate for the common source of price fluctuation is the income and wealth of the very rich. The supply of art is inelastic (at least for deceased artists, who make up the vast majority of our sample), so art prices are determined by demand, which mostly depends on the tastes and wealth of collectors. In contrast to tastes, which are likely to be collector-specific and, thus, unlikely to move aggregate prices, wealth is largely driven by the state of the economy. As Figure 1 highlights, wealth held by the very top earners is also very volatile, which could a priori rationalize why art prices vary so much. Goetzmann et al. (2011) identify top income and wealth as one of the drivers of art prices. They document a long-run relation between top incomes and art prices similar to the relationship between dividends and prices. The leading interpretation in the literature is that top wealth proxies for the unobserved utility dividend a collector would earn from purchasing a diversified art portfolio or, equivalently, the shadow rent the collector would be willing to pay for a collection.

Table 1 displays results for contemporaneous regressions of five-year log art price changes on five-year log-changes in top incomes and top wealth. (We mostly rely on five-year relations throughout the paper to be aligned with the relatively long holding periods in the art market.) We see that art prices are, in fact, not so disconnected from fundamentals: the relationships between art prices and top wealth and incomes are uniformly large and significant and can each explain around one third of changes in aggregate art prices. In contrast, the explanatory power of the U.S. GDP is only about half the one of top wealth. We

Table 1. Art Prices and Wealth

	U.S. top income	U.S. top wealth	UK top income	UK top wealth	U.S. GDP
β_w	0.79***	0.83***	0.68***	0.73***	2.65***
(Standard error)	(0.19)	(0.18)	(0.19)	(0.15)	(0.89)
R^2	0.28	0.34	0.25	0.37	0.18

Notes. This table shows results for contemporaneous regressions of five-year hedonic price log-changes on five-year log changes in several measures of wealth. These measures proxy for the demand for art and correspond to the real per capita income or wealth of the top 1% in the United States or the United Kingdom and GDP growth. The sample period is 1957–2015.

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

conclude that top wealth is a reasonable proxy for the unobserved art dividend.

Do the estimates in Table 1 rationalize why art prices are volatile? Not necessarily. Textbook theory tells us that asset prices react to news about future dividends (e.g., Campbell 1991). The regressions in Table 1, instead, concern contemporaneous changes in dividends. To understand what moves art prices and whether movements in art prices are consistent with unbiased expectations, we need to examine forecasting relations between prices and wealth. In particular, textbook theory predicts that changes in art prices should forecast changes in wealth. In the next section, we show that this standard prediction is not borne out in the data.

3. Predictability of Art Price Changes

The results in the previous section suggest that a common force drives changes in art prices. We also argue that top wealth is a reasonable proxy for the demand for art. In this section, we use lead/lag relations between prices and wealth to ask if art price variation can be justified by future demand.

3.1. Predictive Relations Between Art Prices and Wealth

Textbook theory implies that prices should be forward-looking. As with any other asset, art prices should equal the present value of future dividends, which we argue is best proxied by top wealth. Hence, if investors have unbiased expectations, change in prices should forecast changes in future wealth. To test this hypothesis, we regress past five-year price changes on the log changes in top wealth over the coming five years:[4]

$$\Delta w_{t\to t+5} = 18.30 - \underbrace{0.15}_{(-1.30)} \Delta p_{t-4\to t} + u_{t\to t+5}, \; R^2 = 0.02. \tag{3}$$

In this equation, past price changes $\Delta p_{t-4\to t} \equiv p_t - p_{t-4}$ are observable at time t (i.e., there is no lookahead bias). To the extent wealth proxies for future demand for art, an increase in art prices should serve as a leading indicator of changes in wealth between t and $t+5$. We see that the relationship is not significant and has a counterintuitive sign: price run-ups tend to be followed by declines in top wealth over the next five years. In contrast, Figure 1 (panel (b)) and the result shown in Equation (4) highlight the very strong fit of the opposite relation:

$$\Delta p_{t\to t+5} = 29.02 - \underbrace{1.07}_{(-5.67)} \Delta w_{t-4\to t} + u_{t\to t+5}, \; R^2 = 0.54. \tag{4}$$

The slope is about minus one: a 1% increase in top wealth over the past five years is followed, on average, by a 1% *decrease* in prices over the next five years. This result is at odds with the textbook logic that prices should forecast future fundamentals.

Before turning to the theoretical interpretation of this result, we document its robustness in Table 2: we forecast five-year price index growth in Panel A and actual transaction returns in Panel B. For each transaction in the repeat-sale data set, we collect the price index appreciation over the five-year period prior to the purchase of the artwork. Round-trip returns are then measured as log price changes over the holding period (round-trip), ignoring transaction costs, consistent with Equation (1).[5] Column (1) reproduces the baseline predictive regression of the past five-year change in top U.S. wealth on the next five-year price change; column (2) examines potential price mean reversion and uses the past five-year price change; column (3) forecasts returns using both past wealth change and past price change; column (4) investigates predictability by time period.

We find that the predictive slopes are uniformly negative, so increases in wealth forecast future declines in prices. We also find evidence of price mean reversion (column (2)) although estimates based on the price index (Panel A) are less precise than results based on repeat sales (Panel B). Further, the ability of top wealth to forecast price index changes goes beyond the price mean reversion; the R^2 in the baseline regression is an impressive 0.54 (against 0.11 in the same specification with past five-year price change). Column (3) indicates that top wealth subsumes past price increases when forecasting aggregate prices (Panel A) but not when predicting round-trip returns (Panel B). Column (4) shows that predictability is robust across time periods. We pursue additional robustness checks in Online Appendix C: Tables A.III and A.IV show that predictability is neither sensitive to the choice of the look-back period nor to the choice of the forecasting horizon.

If investors display unbiased expectations with constant risk premiums, price changes should not be forecastable. Our evidence, thus, suggests that either agents have rational expectations and discount rates vary over time or agents have biased expectations. We consider the rational expectations explanations in Section 3.3. Our preferred explanation is that investors over-extrapolate past demand trends. The mechanism works as follows. Suppose the demand for art is negatively autocorrelated: increases in demand are only temporary and fade out over time. Extrapolative investors ignore this feature of the data; they bid prices up when demand increases as if the increase in demand were permanent. When demand eventually weakens, investors are surprised by the (predictable) reversion in the demand, and prices fall.

In the data, we observe that our demand proxy is indeed negatively autocorrelated. When we estimate ϕ as the regression coefficient of $\Delta w_{t\to t+5}$ on $\Delta w_{t-4\to t}$, we find that ϕ is negative at −0.53. The dynamics of

Table 2. Top Wealth Growth Forecasts Art Returns

	(1)	(2)	(3)	(4)
Panel A. Dependent variable: Five-year price change				
Past five-year wealth change	−1.07***		−1.15***	
(Standard error)	(0.22)		(0.30)	
Past five-year price change		−0.36	0.11	
(Standard error)		(0.24)	(0.23)	
Past five-year wealth change (1962–1975)				−0.30
(Standard error)				(1.46)
Past five-year wealth change (1976–1995)				−1.23***
(Standard error)				(0.38)
Past five-year wealth change (1996–2014)				−0.81*
(Standard error)				(0.41)
R^2	0.543	0.106	0.541	0.385
N	49	49	49	49
Panel B. Dependent variable: Round-trip return				
Past five-year wealth change	−0.61***		−0.41***	
(Standard error)	(0.15)		(0.13)	
Past five-year price change		−0.42***	−0.20***	
(Standard error)		(0.08)	(0.02)	
Past five-year wealth change (1962–1975)				−2.30**
(Standard error)				(0.99)
Past five-year wealth change (1976–1995)				−0.88***
(Standard error)				(0.21)
Past five-year wealth change (1996–2014)				−0.53**
(Standard error)				(0.24)
R^2	0.055	0.048	0.060	0.063
N	25,529	25,529	25,529	25,529

Notes. This table shows slopes in a predictive regression of future art prices on past five-year top wealth growth, past five-year hedonic price change, and contemporaneous five-year top wealth growth. In Panel A, the dependent variable is future five-year hedonic log price change; inference is conducted using Hodrick (1992) standard errors for the null of no predictability. In Panel B, the dependent variable is round-trip returns in the sample of repeat-sale transactions, and standard errors are double-clustered at the time of purchase and the name of the artist. Column (1) presents the predictive regression of the past five-year change in wealth on the next five-year/round-trip price change; column (2) uses the past five-year price changes; column (3) forecasts returns using both past wealth change and past price change; column (4) investigates predictability in subsamples. The sample period is 1957–2015.

***$p < 0.01$; **$p < 0.05$; *$p < 0.1$.

prices and wealth, thus, appear to be intertwined, in line with demand over-extrapolation. Over a five-year window, a 1% increase in wealth translates into a contemporaneous similar increase in prices (the point estimate in Table 1 is 0.83). This 1% increase in price mean-reverts in the subsequent five years by about 1% (the point estimate in Table 2 is –1.07). The corresponding coefficient for wealth, ϕ, is also negative, albeit smaller in magnitude. We don't expect the two mean-reversion coefficients to be equal: these simple regressions neglect potential autocorrelation at different frequencies, and mean reversion may occur at different speeds than our five-year window. To get a clearer picture, in Online Appendix D, we study the relation between prices and top incomes in two vector autoregressive models (Tables A.V and A.VI). The models' impulse response functions exhibit mean reversions in similar magnitudes to shocks in both prices and wealth, again consistent with our demand over-extrapolation interpretation.

3.2. Variance Decomposition of Art Price Changes

The results in the previous section are useful to quantify why art prices vary over time. To do so, consider the following regression:

$$\Delta p_{t\to t+5} = \alpha + \beta_w \Delta w_{t\to t+5} + \gamma_w \Delta w_{t-4\to t} + \Delta\tilde{p}_{t\to t+5}, \qquad (5)$$

where this specification captures both the contemporaneous relation between wealth and prices as well as the predictive relation between past wealth and future price changes. The residual term $\Delta\tilde{p}_{t\to t+5}$ captures the part of the price change that cannot be explained by contemporaneous and past changes in wealth. We note in the previous section that wealth is negatively autocorrelated. With some manipulations, we can account for wealth mean reversion and isolate the predictable variation in prices. Denote $\tilde{w}_{t\to t+5}$ as the change in wealth that is unpredictable by past increases in wealth, that is, $\Delta w_{t\to t+5} = \Delta\tilde{w}_{t\to t+5} + \phi\Delta\tilde{w}_{t-4\to t}$. We can rewrite Equation (5) to account for wealth mean-reversion:

$$\Delta p_{t\to t+5} = \alpha + \beta_w \Delta\tilde{w}_{t\to t+5} + (\gamma_w + \beta_w\phi)\Delta w_{t-4\to t} + \Delta\tilde{p}_{t\to t+5}. \qquad (6)$$

We find that the predictive coefficient, adjusted for wealth mean-reversion equals $\gamma_w + \beta_w \phi = -1.10$, which is in the same ballpark as our earlier estimates (of Table 2, Panel A, and Equation (4)).

Ignoring constant terms, the conditional expectation of future price changes is

$$\mathrm{E}_t[\Delta p_{t\to t+5}] = (\gamma_w + \beta_w \phi)\Delta w_{t-4\to t}, \tag{7}$$

where the conditional expectation is taken with respect to Δw_t. Abstracting from the unobserved dividend, $\mathrm{E}_t[\Delta p_{t\to t+5}]$ is an expected return, that is, the discount rate a rational agent would require to invest in a diversified art portfolio at time t. Taking the variance of Equation (6), inserting (7), and dividing by the variance of log price changes gives

$$\begin{aligned} 1 &= \frac{\mathrm{cov}(\Delta p_{t\to t+5}, \Delta p_{t\to t+5})}{\mathrm{var}(\Delta p_{t\to t+5})} \\ &= \frac{\mathrm{cov}(\beta_w \Delta w_{t\to t+5}, \Delta p_{t\to t+5})}{\mathrm{var}(\Delta p_{t\to t+5})} + \frac{\mathrm{cov}(\Delta \tilde{p}_{t\to t+5}, \Delta p_{t\to t+5})}{\mathrm{var}(\Delta p_{t\to t+5})} \\ &\quad + \frac{\mathrm{cov}(\mathrm{E}_t(\Delta p_{t\to t+5}), \Delta p_{t\to t+5})}{\mathrm{var}(\Delta p_{t\to t+5})}. \end{aligned}$$

The first two terms isolate variation resulting from shocks contemporaneous to returns that cannot be forecast by past wealth, and the third term captures the predictable variation in returns. In the spirit of Campbell (1991), we interpret the first two terms as news about future demand and the third one as the discount rate (i.e., expected return) component. The discount rate component reflects the variance of returns that can be attributed to changes in the discount rate a rational agent would require to invest in art. By contrast, we label non–discount-rate shocks as "cash-flow" shocks. The decomposition is agnostic as to whether discount rate changes are consistent with rational expectations, namely whether they reflect changes in a risk premium.

We find that about two thirds (65.6%) of art return variance can be explained by wealth shocks (measured by U.S. top 1% wealth); that is, only 34.4% of the variance cannot be related by changes in wealth. The lion's share of the wealth shocks (59.0%) is attributable to discount rate shocks. Unless this variation can be rationalized by variation of risk premiums, this means that art returns exhibit excess volatility in the sense of Shiller (1981). The variation of art returns would decline by more than half without this predictable variation.[6]

3.3. Rational Explanations?

3.3.1. Can Art Price Predictability Be Explained by Risk? Our analysis indicates that investors could have timed the market using public information available at purchase. However, the fact that changes in prices are forecastable does not necessarily mean that investors have biased beliefs. Public information should not forecast prices unless it correlates with a time-varying discount rate. For example, past wealth increases could forecast lower returns because these increases correlate with economic booms and a lower risk aversion as in the model of Campbell and Viceira (1999). Thus, before turning to a behavioral interpretation, we ask if risk could rationalize our predictability results. Taking expectations in Equation (4) and subtracting the real risk-free rate gives us, under rational expectations, the (log) risk premium required by diversified art investors.[7] Risk could, therefore, rationalize our findings provided the art risk premium correlates with lagged wealth growth. Still, the risk premium would then have to be extremely volatile. We already note that about half of aggregate price fluctuation should be explained by variation in this risk premium. For instance, given that the volatility of five-year wealth changes amounts to 26%, the coefficient in column (1) of Table 2 implies a five-year premium that varies by $26 \times -1.07\% \approx 28\%$ per year, which seems unrealistically high.

The risk premium interpretation is also more challenging when return forecasts are lower than the risk-free rate or negative (Schwert 2003, Fama 2014). We find that expectations of future price changes are often strongly negative. This is readily apparent from Figure 1 for all art sales, but we find similar results when looking at round-trip returns. To see this, we use transactions in the repeat-sale sample and consider strategies that buy all available artworks when top wealth has increased over a five-year time frame. These strategies are constructed out-of-sample and, thus, mimic the returns of a hypothetical (perfectly diversified) investor in real time. In untabulated results, we find that such strategies frequently earn negative returns. For instance, a strategy that buys artworks when top wealth has increased over a five-year time frame by more than 40% (which corresponds to the top 25 percentile) earns an average round-trip return of −14.9%. This constitutes a lower bound because we work with raw returns rather than excess returns, and we ignore transaction costs. This is much stronger evidence of negative return predictability than in other asset markets. For instance, Greenwood et al. (2019) find that betting against stock market bubbles is extremely hard: a price run-up of 100% of an industry rarely leads to strong sudden negative price adjustments even when a longer time window of two subsequent years is considered.

Furthermore, asset pricing theory imposes tight conditions on the behavior of the risk premium, which must depend on the amount of risk faced by investors as well as their willingness to bear risk. In the spirit of Greenwood and Hanson (2015), we, thus, examine under what conditions risk could rationalize our findings. Theory predicts that the risk premium on any asset returns depends on the covariance between future returns and the marginal utility of diversified investors. In our context, this means that

$$\mathrm{E}_t[\Delta p_{i,t\to t+h}] + \frac{1}{2}\mathrm{var}_t[\Delta p_{i,t\to t+h}] - r_{f,t\to t+h} = \mathrm{cov}_t[-m_{t\to t+h}, \Delta p_{i,t\to t+h}] \\ = \mathrm{corr}_t[-m_{t\to t+h}, \Delta p_{i,t\to t+h}]\mathrm{var}_t[m_{t\to t+h}]\mathrm{var}_t[\Delta p_{i,t\to t+h}], \tag{9}$$

where $m_{t\to t+h}$ denotes the log of the stochastic discount factor over the period h and $r_{f,t\to t+h}$ is the log risk-free rate. The left-hand side of Equation (9) equals the log risk premium on art investments, adjusted for Jensen's inequality by adding one half of the variance of art's log returns. The right-hand side of this equation implies that variation in the art risk premium must either result from a time-varying correlation between art returns and investor utility $\mathrm{corr}_t[-m_{t\to t+h}, \Delta p_{i,t\to t+h}]$, variation in the economy-wide price of risk $\mathrm{var}_t[m_{t\to t+h}]$, or variation in the risk of art investments $\mathrm{var}_t[\Delta p_{i,t\to t+h}]$.

The first term of Equation (9), $\mathrm{corr}_t[-m_{t\to t+h}, \Delta p_{i,t\to t+h}]$, captures art's hedge value to investors. High realizations of $-m_t$ correspond to states of the world in which the marginal utility of diversified investors is low, that is, good states of the world for art investors. A positive correlation implies that art has relatively higher returns in good states of the world. Art has a negative hedge value for diversified investors and must, therefore, command a positive risk premium, everything else equal. To explain our predictability result, this correlation must vary substantially over time, which requires a corresponding time-varying hedge. Furthermore, the correlation would have to be low when expected returns are high, that is, when wealth has been increasing over time. This seems highly unlikely.

Turning to the second term of Equation (9), a more plausible explanation is that the art risk premium varies because the premium on all risky assets $\mathrm{var}_t[m_{t\to t+h}]$ varies over time. The literature proposes three main mechanisms for why prices of risk may vary over time: habit formation models in which investor risk aversion varies over time as in Campbell and Viceira (1999), long-run risk models in which the long-run growth rate varies over time as in Bansal and Yaron (2004), and rare disaster models in which macroeconomic tail risk varies over time (Gabaix 2012, Wachter 2013). We, therefore, test whether the predictable variation in art prices can be explained by omitted economy-wide prices of risk. We present results in Table 3, in which we add variables that have forecasting power for equities and bonds, namely the log dividend yield, the term spread, and the default yield spread, obtained from Goyal and Welch (2008).[8] Adding these risk premium proxies to the basic Model 1, one at a time (Models(2–4), and all at once (Model 5), yields that the ability of past wealth growth to predict future price changes remains essentially unchanged. The three additional risk premiums do not predict art returns (neither aggregate returns, Panel A, nor round-trip returns, Panel B).

Finally, the third term of Equation (9) suggests another reasonable explanation, namely that the riskiness of art investments $\mathrm{var}_t[\Delta p_{i,t\to t+h}]$ is higher when top wealth has been growing over time. We, thus, test if top wealth forecasts the volatility of future returns. Our realized volatility proxy $\varsigma_{t\to t+5}$ consists of the five-year change in hedonic price dispersion (in the time series setting) and of the absolute round-trip returns (in the repeat-sale setting). We find that wealth growth is not statistically significantly related to future volatility. For instance, in the time series setting, we find

$$\Delta\varsigma_{t\to t+5} = 0.08 - \underbrace{0.07}_{(-0.83)}\Delta w_{t-4\to t} + u_{t\to t+5},\ R^2 = 0.04. \tag{10}$$

Likewise, wealth growth has no forecasting power for absolute or squared round-trip returns. We conclude we cannot reconcile our predictability results with rational expectation theories in which risk premiums vary over time.

3.3.2. Credit Frictions Explanation. Even if art investors have rational expectations, credit frictions could make art returns predictable. Although art investors rarely borrow to buy art, aggregate credit frictions could still play an indirect role affecting the art market. For instance, when credit is scarce, the opportunity cost of holding art increases, such that investors require a higher expected return not to sell their art. This mechanism could explain why art returns are higher in downturns; namely this would be the case if negative wealth growth correlates with more binding credit constraints.

To address that possibility, we obtain data on household debt from Jordà et al. (2017). The data cover 17 advanced economies since 1870 on an annual basis. We measure debt levels using the total loans to households over GDP and mortgage loans to the nonfinancial private sector over GDP. In Table 4, we forecast art returns using five-year changes in these two debt-to-GDP ratios. Specifically, as in Table 3, we maintain our baseline specification with wealth changes on the right-hand side and ask if the predictability disappears once we control for credit cycle proxies. Panel A uses a time series regression in which we use U.S. data to forecast aggregate price changes. In this specification, we find that U.S. credit is not statistically significantly related to future art returns. In Panel B, we forecast transaction returns, which allows us to use country-level credit data although wealth changes are still measured at the U.S. level. Here, we find that credit changes forecast art returns with a

Table 3. Predictability Result Controlling for Traditional Predictors

	(1)	(2)	(3)	(4)	(5)
Panel A. Dependent variable: Five-year price change					
Past five-year wealth change	-1.07^{***}	-1.12^{***}	-1.08^{***}	-1.07^{***}	-1.16^{***}
(Standard error)	(0.22)	(0.21)	(0.23)	(0.22)	(0.24)
Past dividend yield		−4.15			−0.16
(Standard error)		(6.19)			(0.20)
Past term spread			−1.07		−2.40
(Standard error)			(3.26)		(3.35)
Past default yield spread				−0.87	5.77
(Standard error)				(9.98)	(11.21)
R^2	0.543	0.550	0.536	0.534	0.538
N	49	49	49	49	49
Panel B. Dependent variable: Round-trip return					
Past five-year wealth change	-0.61^{***}	-0.61^{***}	-0.61^{***}	-0.60^{***}	-0.63^{***}
(Standard error)	(0.15)	(0.18)	(0.13)	(0.15)	(0.16)
Past dividend yield		−0.66			7.07^{*}
(Standard error)		(3.96)			(3.93)
Past term spread			−0.18		−0.09
(Standard error)			(1.89)		(0.12)
Past default yield spread				2.82	−0.79
(Standard error)				(2.25)	(1.44)
R^2	0.055	0.055	0.055	0.055	0.057
N	25,529	25,529	25,529	25,529	25,529

Notes. This table shows slopes in a predictive regression of future art prices. In Panel A, the dependent variable is future five-year hedonic log price change; inference is conducted using Hodrick (1992) standard errors for the null of no predictability. In Panel B, the dependent variable is round-trip return in the sample of repeat-sale transactions, and standard errors are double-clustered at the time of purchase and the name of the artist. The first column reproduces the baseline result of art price predictability by wealth growth shown in column (1) of the respective panels of Table 2. We expand these results in columns (2)–(5) by including proxies for economy-wide risk premiums on equities and bonds, namely the dividend yield (the difference between the log of dividends and the log of lagged prices), the term spread (the difference between long term yield on government bonds and treasury-bill rate), and the default yield spread (the difference between Baa and Aaa-rated corporate bond yields). The sample period is 1957–2015.

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

negative sign (Models 2 and 4), meaning that periods of high credit growth are followed by lower returns. We also find that the forecastability by wealth growth remains unchanged.

The fact that credit variables forecast lower returns is consistent with recent evidence showing that credit booms tend to be followed by lower real estate prices and lower stock returns (Di Maggio 2017, Davis and Taylor 2019). However, this result does not necessarily imply that speculators are constrained in bad times. Credit booms tend to be followed by slower growth, deeper recessions, and increased likelihood of financial crises (e.g., Mian and Sufi 2009, Schularick and Taylor 2012, Baron and Xiong 2017, Mian et al. 2017). Instead of time-varying credit frictions, the literature tenet emphasizes excessive credit as the causal mechanism for this body of results. To distinguish the two potential mechanisms, we separate positive from negative credit growth rates in columns (3) and (5) of Table 4. We find that positive credit growth is more strongly associated with future returns, and negative growth has no forecasting power. This is consistent with the general interpretation in the literature that credit booms are periods of excessive credit. We conclude that credit frictions are unlikely to explain art prices' forecastability by past changes in wealth.

To summarize, our evidence indicates that art prices often deviate from fundamental value. Prices increase with good news about top income and wealth, but they rise *too much*, so that increases in wealth are followed by declines in prices. We argue that rational expectations, in the form of time-varying risk premiums or credit frictions, are unlikely to explain the evidence. Our preferred interpretation is that the marginal investor sometimes exhibits too optimistic expectations. Models of cash flow (i.e., demand) extrapolation, in particular, predict that past growth in fundamentals, rather than past prices, forecast future returns negatively as we find in the data. Not all investors need to display incorrect beliefs because short-sale constraints prevent prices from reflecting the opinion of less optimistic investors (Miller 1977). In turn, if investors have diverse opinions, we expect them to trade more during booms as we document in the next section.

4. Properties of Art Price Booms and Busts

We argue that some investors exhibit extrapolative beliefs, pushing prices away from demand fundamentals, leading to predictable price changes. This reasoning begs the question of why rational investors don't exploit these potential profit opportunities. One potential explanation is that it is essentially impossible to sell art short. When prices exceed demand fundamentals, a rational investor has no better option than to sell to an optimist. This explanation predicts that prices and volume are positively correlated and that investors trade because they disagree; volatile disagreement, in turn, can generate price bubbles as investors trade aggressively against one another (Scheinkman and Xiong 2003). The presence of extrapolative investors can further rationalize why investors exhibit differences of opinion in the first place (Barberis et al. 2018). In this section, we document new facts relating prices, volume, and disagreement to inform these theories.

4.1. Trading Volume

We note in the introduction that art prices tend to coincide with high trading volume (Figure 1, panel (c)).

Table 4. Predictability Result Controlling for Credit Growth Predictors

	(1)	(2)	(3)	(4)	(5)
Panel A. Dependent variable: Five-year price change					
Past five-year wealth change	−1.07***	−1.08***	−1.09***	−1.12***	−1.14***
(Standard error)	(0.22)	(0.23)	(0.23)	(0.24)	(0.23)
Past five-year credit change		0.29			
(Standard error)		(1.88)			
Past five-year credit change (+)			1.99		
(Standard error)			(2.49)		
Past five-year credit change (−)			−2.89		
(Standard error)			(3.87)		
Past five-year mortgage change				0.86	
(Standard error)				(1.71)	
Past five-year mortgage change (+)					1.63
(Standard error)					(2.27)
Past five-year mortgage change (−)					−0.76
(Standard error)					(3.87)
R^2	0.543	0.534	0.542	0.542	0.537
N	49	49	49	49	49
Panel B. Dependent variable: Round-trip return					
Past five-year wealth change	−0.61***	−0.60***	−0.60***	−0.57***	−0.57***
(Standard error)	(0.15)	(0.14)	(0.14)	(0.15)	(0.15)
Past five-year credit change		−0.35***			
(Standard error)		(0.09)			
Past five-year credit change (+)			−0.43***		
(Standard error)			(0.04)		
Past five-year credit change (−)			0.01		
(Standard error)			(0.51)		
Past five-year mortgage change				−0.79**	
(Standard error)				(0.35)	
Past five-year mortgage change (+)					−1.01
(Standard error)					(0.64)
Past five-year mortgage change (−)					0.38
(Standard error)					(1.13)
R^2	0.055	0.057	0.057	0.059	0.059
N	25,529	25,186	25,186	25,529	25,529

Notes. This table shows slopes in a predictive regression of future art prices. In Panel A, the dependent variable is future five-year hedonic log price change; inference is conducted using Hodrick (1992) standard errors for the null of no predictability. In Panel B, the dependent variable is round-trip return in the sample of repeat-sale transactions, and standard errors are double-clustered at the time of purchase and the name of the artist. The first column reproduces the baseline result of art price predictability by wealth growth shown in column (1) of the respective panels of Table 2. We expand these results in columns (2)–(5) by including data on household debt from Jordà et al. (2017). In Panel A, we use the total loans to households over GDP for the United States. In Panel B, we use the mortgage loans to the (nonfinancial) private sector over GDP for 17 advanced economies. The sample period is 1957–2015.

***$p < 0.01$; **$p < 0.05$; *$p < 0.1$.

The correlation between annual log-changes in prices and log-changes in volume is 0.38 in our data. This positive correlation is not specific to the art market. Fluctuations in transaction volume also accompany price cycles in the stock and real estate markets (e.g., Stein 1995, Genesove and Mayer 1997, Hong and Stein 2007). Greenwood et al. (2019), in particular, finds that in U.S. industry returns, attributes of price run-ups, including turnover, can help forecast an eventual bust. However, other authors propose different explanations for this correlation between prices and volume, including down payment effects (Stein 1995), sellers' loss aversion (Genesove and Mayer 2001), and differences of opinion (Hong and Stein 2007).[9]

To better understand why trading volume fluctuates, we decompose trading volume by holding period. It seems plausible that, in downturns, the number of consignments disproportionally represents forced sales. In that case, short-term speculators may step in to purchase artworks at a discount, meaning that speculators play a stabilizing force. This channel predicts that short-term transactions become more prevalent in downturns. Inversely, many models predict that speculators have a destabilizing effect on prices either because they exhibit biased expectations (e.g., Defusco et al. 2020) or they rationally decide to ride bubbles (e.g., Abreu and Brunnermeier 2003). Either way, these models predict that short-term transactions become more prevalent during booms. Figure 2 shows

Figure 2. (Color online) Volume Decomposition: The Role of Short-Term Volume

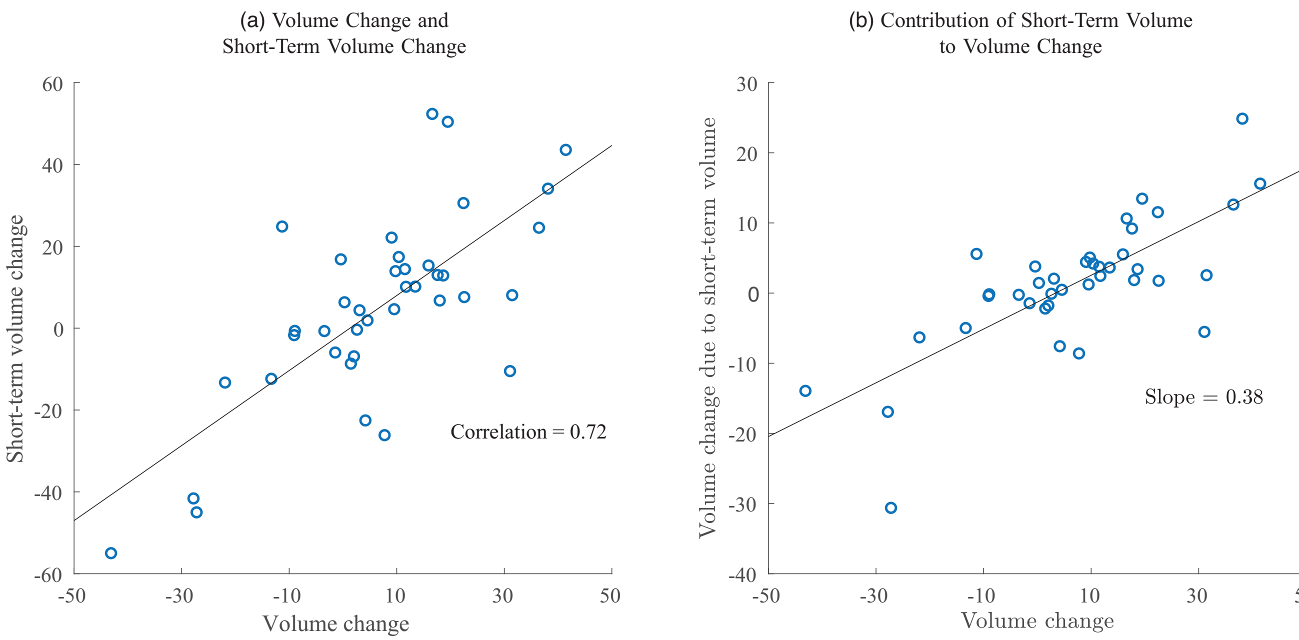

Notes. (a) Volume change and short-term volume change. (b) Contribution of short-term volume to volume change. This figure quantifies the relation between volume and short-term transactions. Short-term transactions are defined as repeat-sale transactions with holding periods less than two years. Both panels are scatterplots with volume change on the *x*-axis. Volume is defined as the yearly number of transactions. Panel (a) shows the yearly percentage change in short-term transactions on the *y*-axis. Panel (b) quantifies the importance of short-term transactions as a component of volume by showing volume change against volume change resulting from short-term transactions. The sample period is 1976–2015.

that a substantial share of variation in trading volume is due to short-term transactions, defined as repeat-sale transactions with holding periods of two years or less.

Panel (a) of Figure 2 relates changes in global resale volume to changes in the number of short-term transactions. The number of short-term transactions is strongly procyclical (the correlation with resale volume is 0.72) and is about twice more volatile. In panel (b), we ask how much of the percentage change in resale volume is due to short-term transactions, a methodology we borrow from Defusco et al. (2020). On the *x*-axis, we again show the change in total resale volume, and on the *y*-axis, we report the variation of total volume resulting from short-term transactions (i.e., the change in the number of short-term trades between t and $t + 1$ divided by time-t total volume). We also plot the least-squares line: the slope indicates that the elasticity of total trading volume to short-term transactions is 0.38. This result is consistent with flippers contributing to the art market's cyclicality. Nevertheless, it could be that rising prices cause holding periods to be shorter, for example, if collectors are more likely to sell after gains than after losses. To address this concern, we repeat our analysis focusing on the subsample of transactions that ended up in nominal losses. We find that the number of losing short-term transactions also goes up with the total number of losing transactions; the correlation with global volume is about unchanged (it is 0.71 in this subsample versus 0.72 in the full repeat-sale sample). Further, this result is similar, both economically and quantitatively, to the findings in Bayer et al. (2020) and Defusco et al. (2020) in that short-term transactions explain a large part of changes in trading volume during the U.S. real estate bubble.

Newspaper articles also suggest that the most glamorous styles (postwar and contemporary art) are more likely to be traded in booming markets. We hypothesize that speculative trading may concentrate on more glamorous art, such as emerging artists' work.[10] If this were true, we expect to see glamour artists being traded more in booms than in busts. Figure 3 shows that glamour art accounts for a significant fraction of total volume growth. We classify postwar art (i.e., art classified as Abstract Expressionism, Pop Art, and modern and contemporary art) as glamour art. Over our sample period, the share of postwar art rises from zero to about 25% over time. We again focus on the 1976–2015 period, in which postwar art always amounts to at least 5% of yearly trading volume. In panel (a), we compare changes in total volume to changes in postwar art volume. In panel (b), we quantify the fraction of the total volume changes resulting

Management Science, 2022, vol. 68, no. 7, pp. 4939–4963, © 2021 The Author(s)

Figure 3. (Color online) Volume Decomposition: The Role of Postwar Art Volume

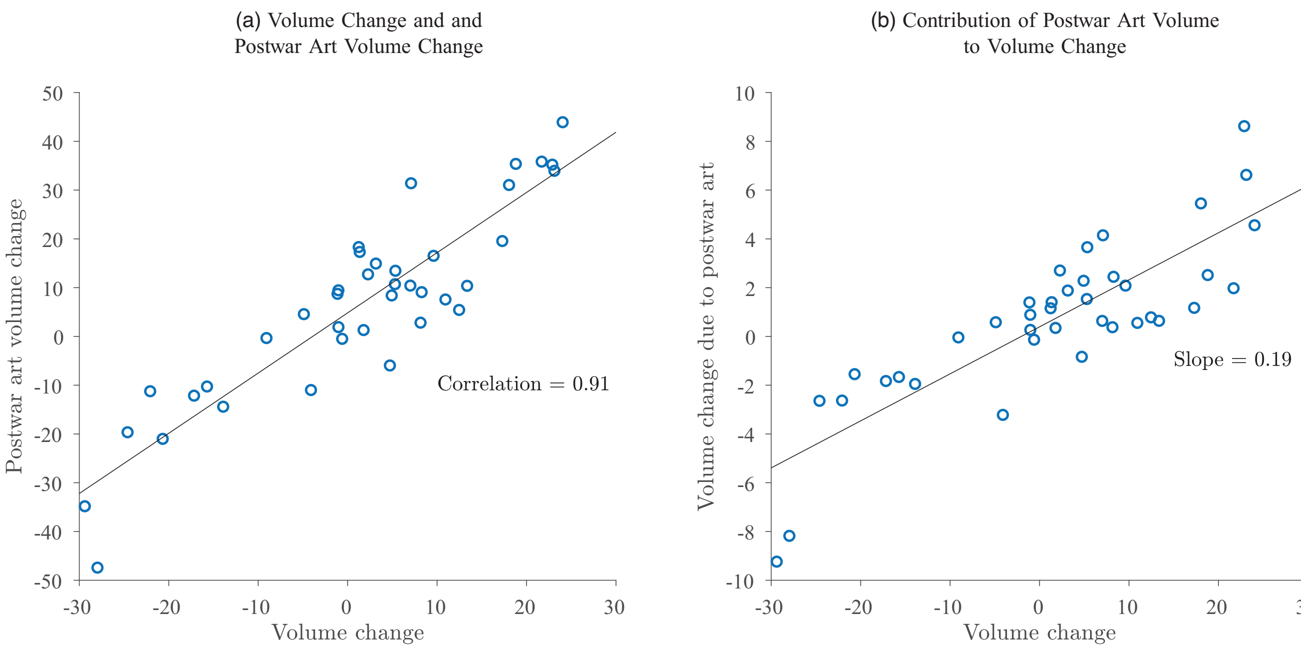

Notes. (a) Volume change and and postwar art volume change. (b) Contribution of postwar art volume to volume change. This figure quantifies the relation between volume and postwar art transactions. Postwar artists are artists classified as Abstract Expressionism, Pop Art, or modern and contemporary art. Both panels are scatterplots with volume change on the *x*-axis. Volume is defined as the yearly number of transactions. Panel (a) shows the yearly percentage change in postwar art transactions on the *y*-axis. Panel (b) quantifies the importance of postwar art transactions as a component of volume by showing volume change against volume change resulting from postwar art transactions. The sample period is 1976–2015.

from postwar art. We find that about 19% of volume growth is due to postwar art volume growth. This suggests that more speculative art is traded more frequently in times of growing art markets.

4.2. Price Dispersion

Another outstanding feature of the art market is that auction prices tend to be more dispersed during booms. Figure 4 shows a scatterplot of percentage changes of prices and in price dispersion (a proxy for volatility). We measure price dispersion by computing, every year, the cross-sectional standard deviation of the residuals in the hedonic regression (1), which captures the yearly price dispersion that cannot be attributed to observed characteristics. Figure 4 depicts that price dispersion exhibits substantial variation over the sample period (1957–2015) and correlates with art prices (the correlation is 0.33).

4.3. Belief Dispersion: Survey Evidence

Our evidence suggests that periods of excessive valuations coincide with intense speculative trading. In speculative trading models, investors trade because they disagree about the value of the traded asset. A candidate reason for investors to disagree is that some over-extrapolate, and others do not (Barberis et al. 2018). In what follows, we use survey data to ask if disagreement correlates with prices, volume, and price dispersion. Subsequently, we ask if belief dynamics are consistent with extrapolative beliefs.

Our data are provided by ArtTactic and consist of price expectations from a pool of art market participants, previously used in Pénasse et al. (2014). Surveyed investors are asked about their financial outlook of a sample of artists over a six-month horizon. The data are sampled semiannually from November 2005 to November 2012 and cover 21 American and European modern and contemporary artists.[11] Investors are given a threefold possible answer ("positive," "neutral," or "negative"), which we code as $+1, 0, -1$. The average belief, thus, corresponds to the difference between positive and negative answers, which is standard in the literature (e.g., Greenwood and Shleifer 2014). We define belief dispersion as the standard deviation of the answers.

We first test if belief dispersion comoves with prices, volume, and price volatility. Panel A in Table 5 reports the results of a regression of the log changes in belief dispersion on the log changes in aggregate prices, volume, and price dispersion. In agreement with speculative trading models, we find that belief dispersion is positively correlated with changes in price, volume, and price dispersion although the regression coefficient is only statistically significant for volume.

Figure 4. (Color online) Prices Changes and Changes in Price Dispersion

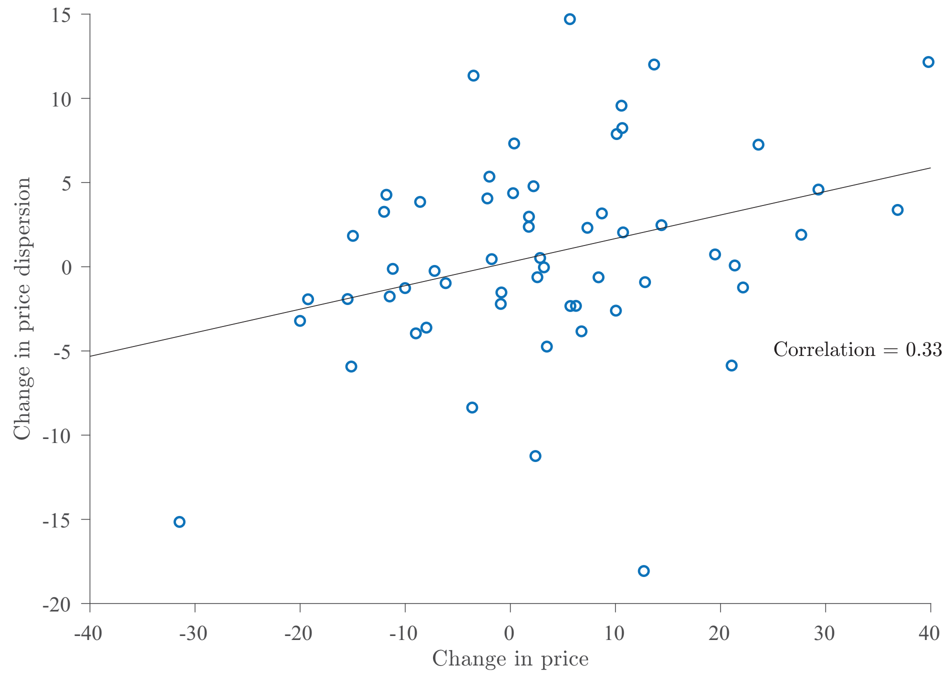

Notes. Scatterplot of annual percentage changes in aggregate art prices against percentage changes in the dispersion of art prices over the period 1957–2015. The art price index is constructed by means of the hedonic regression described in Equation (1). Price dispersion is the cross-sectional standard deviation of the residuals in this hedonic regression and captures the yearly price dispersion that cannot be attributed to observed characteristics.

We next examine in Panel B of Table 5 the relation between average beliefs and past increases in wealth and art prices. In Table 2, we show that past increases in wealth and prices forecast negative returns. Hence, if surveyed investors hold unbiased beliefs, our measure of average beliefs should be negatively correlated with past prices and wealth. In contrast, if the average surveyed investor displays extrapolative beliefs, we expect positive relations. Given that the ArtTactic data has a very short time coverage (2005–2012) and measures investors' expectations over a short horizon, we regress investors' expectations on one-year changes in wealth and prices, whereas Table 2 uses five-year changes. We rescale the expectation series in Table 5 to have the same standard deviation as realized artist returns, which we obtained from Tutela Capital. This ensures that the expected returns series are of similar magnitude as realized returns series.

The first column of Panel B indicates that, when wealth has been growing, investors expect higher returns going forward, a statistically significant relation. We remark that the coefficient is smaller than the coefficient we observe for future realized returns in Table 2. Although this can be explained by the different data sources, we remark that the coefficient should be smaller if some—but not all—surveyed investors extrapolate. If some investors do not extrapolate, they would like to sell short when they believe prices exceed fundamentals. However, because selling art short is essentially impossible, their beliefs may not be reflected in prices and at the same time be included in our survey data. As a result, the survey responses of investors who hold correct beliefs would push down the coefficient in Panel B, which may explain why the coefficients are smaller than in Table 2.

The second column of Panel B of Table 5 shows the regression results of average beliefs on past returns. We see that, when recent past returns are high, investors expect higher future returns. This is consistent with the stock market evidence in Greenwood and Shleifer (2014). The coefficient is smaller than the coefficients associated with past wealth and becomes even smaller when controlling for past wealth (fourth column). This is in line with our results in Table 2, which we interpret as investors extrapolating past fundamentals rather than past returns. In the remaining columns, we introduce artist-level returns (i.e., the log change in the artist price index). These series allow us to distinguish the relationship between beliefs and past returns measured at the artist and aggregate levels. We observe that investors' expectations are less sensitive to past artist returns than to past market returns. In untabulated regressions, we find that our results are unchanged when we include controls for variables that may affect expected returns, the log dividend yield, the term spread, and the default yield spread (as in Table 3).

Table 5. Survey Evidence

Panel A. Changes in belief dispersion				
	(1)		(2)	(3)
Price changes	7.36			
(Standard error)	(6.34)			
Changes in volume			7.15***	
(Standard error)			(1.71)	
Changes in price dispersion				34.38
(Standard error)				(23.73)
R^2	0.007		0.027	0.019
N	209		209	209
Panel B. Determinants of investors' beliefs				
	(1)	(2)	(3)	(4)
Past wealth change	0.52***			0.43*
(Standard error)	(0.15)			(0.22)
Past market return		0.38*		0.05
(Standard error)		(0.19)		(0.15)
Past artist return			0.15*	0.13
(Standard error)			(0.09)	(0.09)
R^2	0.048	0.018	0.021	0.058
N	230	230	230	230
Panel C. Determinants of belief dispersion				
	(1)	(2)	(3)	(4)
Past wealth change	2.21***			1.78***
(Standard error)	(0.81)			(0.37)
Past market return		2.02**		1.24***
(Standard error)		(0.83)		(0.46)
Past artist return			0.66*	0.43
(Standard error)			(0.34)	(0.31)
R^2	0.022	0.024	0.017	0.044
N	230	230	230	230

Notes. Panel A shows the regression results of the log of changes in belief dispersion on the log of changes in art prices, log of changes in volume, and log of changes in price dispersion. Panel B shows regressions of average beliefs on past one-year top wealth growth, past aggregate market returns, and past artist returns. Panel C shows regressions of belief dispersion on the same variables. Survey expectations are measured using ArtTactic survey data on the short-term confidence in the market for each artist. Artist price indexes are provided by Tutela Capital. Market price is the cross-artist average price. Variables are sampled at the semiannual frequency and cover at most the period 2006:S1–2012:S2. Standard errors are clustered at the period level.

***$p < 0.01$; **$p < 0.05$; *$p < 0.1$.

Panel C reports the regression results with dependent variable belief dispersion measured in levels and the same independent variables. Belief dispersion series are standardized to have zero mean and unit variance. We observe that disagreement is positively related to wealth growth (Specification (1)), market returns (Specification (2)), and artist returns (Specification (3))—with the former dominating the latter (Specification (4)). The effect of wealth and market returns is fairly large. For instance, a one standard deviation (7.3%) increase in top wealth increases disagreement by $1.8 \times 0.073 \approx 13\%$. Artist lagged return plays a minor role in comparison, consistent with the results in Panel B. Overall, the evidence is consistent with extrapolative beliefs contributing significantly to disagreement.

4.4. A Case Study: The Late 1980s Bubble

Although we emphasize that our predictability results hold throughout our sample, it is interesting to zoom in on the period 1985–1995 as this period (which embeds the late 1980s bubble) corresponds to the largest boom–bust in the art market over the postwar era. Many commentators at the time stressed the role of Japanese collectors in the development of the art bubble. Indeed, the timing of the art market boom–bust coincides with the timing of the Japanese asset price bubble. Hiraki et al. (2009) document that art imports into Japan were strongly positively correlated with both art prices and Japanese stock prices during the Japanese "bubble period." Although this narrative is interesting and credible, it is unlikely that the influence of Japanese markets alone was sufficient to explain the dramatic market fluctuations over the period. Hiraki et al. (2009) show that the Japanese buying spree was mostly concentrated on French Impressionist and Post-Impressionist art, which is particularly attractive to Japanese culture, notably because of the influence of Japanese woodblock print techniques and themes on Impressionist art. However, the boom and bust were not limited to Impressionist Art: Pénasse et al. (2014) show that prices increased, in different magnitudes, across all market segments. The price increases over 1985–1990 almost perfectly predicted the subsequent declines over 1990–1995.

Figure 5 summarizes and complements these prior results. Panel (a) shows price indices and volume for Impressionist Art, postwar art, and other market segments.[12] Panel (b) plots trading volume, detrended by taking the five-year log-change as in the earlier sections, against average future 12-month price changes as in Barberis et al. (2018). Panel (c) compares aggregate volume and price levels. We observe that Impressionist Art prices quadrupled over the 1985–1990 period although volume increased by about 50%. However, market growth was not limited to Impressionist Art. In fact, postwar art experienced an even stronger increase in both prices and volume in line with our results in Section 4.1. This pattern is reminiscent of the explosion of high-tech stock prices and volume during the internet bubble of 1998–2000 (see Hong and Stein 2007, figure 1).

Historical accounts of financial bubbles suggest that bubbles develop in a context of radical innovation or fundamental uncertainty. Our findings suggest that the entry of Japanese investors is likely to have acted as a precipitating event to borrow Shiller's (2000) terminology. This narrative suggests that this event attracted speculators who pushed prices and volume further up (the share of short-term transactions rose 47% over the 1985–1989 period). Barberis et al. (2018), in particular, propose a model in which endogenous disagreement arises because some investors have

Figure 5. (Color online) The Late 1980s Bubble

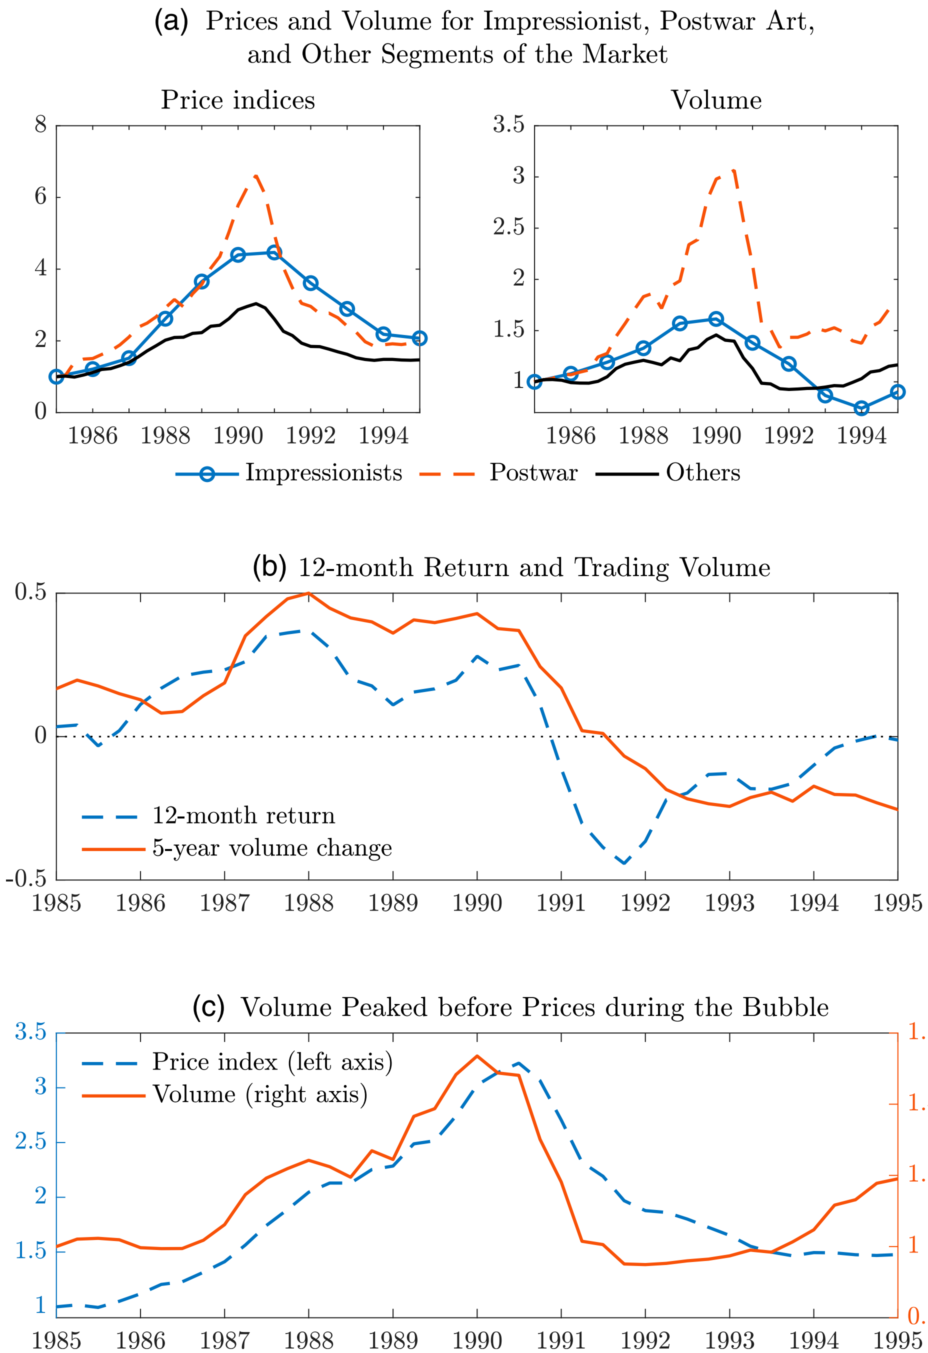

extrapolative beliefs. A series of positive cash flow news increases prices, which attracts extrapolators who trade with fundamental traders. Extrapolation exacerbates disagreement that grows endogenously over the course of the bubble. Barberis et al. (2018) suggest comparing the correlation between past returns and volume as a diagnostic of their model. Using data from several historical bubbles, they show that this correlation is systematically larger during the bubble period than in the aftermath of the bubble. We conduct an out-of-sample test of their theory in panel (b) of Figure 5, which shows the past 12-month return (log price change) and trading volume (detrended by taking a five-year log change) similarly to figure 6 in Barberis et al. (2018). We observe a strong relationship between volume and the past 12-month return during the bubble period, which disappears after 1991. The correlation between the two series is 0.78 between January 1985 and December 1991. In contrast, the correlation in the three-year postbubble period from January 1992 to December 1995 is merely −0.02. Hence, the late 1980s bubble does not seem different from other historical episodes and is consistent with the prediction of Barberis et al. (2018).

Also noteworthy is that volume reached its maximum a year before the peak in aggregate prices (panel (c) of Figure 5). Hong and Stein (2007) and Defusco et al. (2020) document similar facts in the internet and U.S. real estate bubble, respectively. Defusco et al. (2020) show how departing from Walrasian market clearing can reproduce this feature of the data when agents exhibit extrapolative expectations. In their model, optimist sellers overestimate the future growth in demand and, thus, set high listing prices. This causes trading volume to fall although the prices of the transactions that actually occur remain high.

This mechanism could be at work in the art market. When putting a work for sale, a seller sets a secret reserve price below which the seller is not willing to depart from the work. Thus, we can verify if the sales rate fell before prices as volume falls, reflecting the fact that sellers had overoptimistic expectations at the peak of the bubble. As our data set only includes items that actually sold, we construct a proxy for the sales rate. For each auction, we divide the yearly number of observed transactions by the maximum lot number. Then we take, for each year, the average sales rate across auctions as our proxy for the aggregate sales rate. This gives us a noisy sales rate proxy, but we do find that, as with volume, the sales rate falls before prices. The decline is about 5% points from a peak of about 23% in 1990 to a trough of roughly 18% in 1993. This is quite a large decline, which is very much in line with the Defusco et al. (2020) explanation.

5. Dissecting Short-Term Transactions

In Section 4.1, we document the procyclical nature of short-term transactions. We argue that short-term transactions primarily reflect the trades of flippers. Art flippers purchase art with the intent of offering the piece at a drastically higher price shortly after purchase. Flipping seems a risky strategy given the high transaction costs in the art market, and yet there is plenty of evidence that such flips sometimes earn very high returns. In one striking example, Jean-Michel Basquiat's *Warrior* sold three times at auction between 2005 and 2012, the painting's price soaring during those seven years by 450% to nearly $9 million. A popular strategy consists in purchasing works of art in galleries to quickly resell them for a profit at auctions, but this practice is also common in the secondary market, as the preceding example suggests. In this section, we try to answer whether speculators are sophisticated and whether they play a correcting force during booms. We first ask what the determinants are of holding periods and subsequently focus on the relative performance of short-term transactions.

5.1. Determinants of Holding Periods

Table 6 shows regression results for the holding period as a function of ex ante characteristics. As a holding period is mechanically affected by the year of purchase, all regressions include year fixed effects as well as a control variable capturing whether the artist died during the holding period. (Our results are not sensitive to this control with two exceptions we highlight as follows.) All characteristics are available at purchase, meaning that we are trying to forecast holding periods in the cross-section. We consider the following variables: price-estimate is the log real hammer price minus the log of the mid of the presale estimates (high versus low estimates) at purchase; log real hammer is the purchase price; major auction house is a dummy equal to one if the artwork was purchased in one of Sotheby's or Christie's London or New York branches; an artist' fame is measured as the number of times the artist is mentioned in Google's digitized books; living artist is an indicator that is equal to one if the artist is alive at the purchase date; a postwar dummy equals one if the artist's style is classified as Abstract Expressionism, Pop Art, or modern and contemporary art. Right-hand side variables are standardized to have a zero mean and unit variance.

Table 6 presents the results. It seems plausible that investors are more likely to overbid for artworks that represent an emotional investment. We observe that the higher the prices are relative to the auctioneer's estimate, the longer the holding period. This is consistent with our interpretation that short-term transactions are more likely to be speculative. We also

Table 6. Determinants of Holding Periods

	Coefficient	(Standard error)	R^2	N
Price estimate	0.47***	0.04	0.09	4,695
Log real price	0.88***	0.07	0.08	25,733
Major auction house	0.83***	0.07	0.08	25,733
Fame	0.39***	0.05	0.07	23,238
Living artist	−0.84***	0.05	0.07	25,733
Postwar artist	−0.09**	0.04	0.07	25,733

Notes. This table shows coefficients for regressions of the holding period on transaction and artist variables. All variables are available at the time of purchase. The regressions include year fixed effects and a control variable capturing whether the artist passed away over the holding period (the dummy then equals one). We consider the following variables: price estimate is the log real hammer price minus the log of the mid of the presale estimates (high versus low estimates) at purchase; log real hammer is the purchase price; major auction house is a dummy equal to one if the artwork was purchased in one of Sotheby's or Christie's London or New York branches; an artist's fame is measured as the number of times the artist is mentioned in Google's digitized books; living artist is an indicator that is equal to one if the artist is alive at the purchase date; a postwar dummy equals one if the artist's style is classified as Abstract Expressionism, Pop Art, or modern and contemporary art. The sample period is 1957–2015.

***$p < 0.01$; **$p < 0.05$; *$p < 0.1$.

find that more expensive artworks, objects from more famous artists, and artworks sold in more prominent auction houses are held for a longer period. This is consistent with speculators focusing on the less scrutinized segments of the market or speculators being less affluent. Finally, holding periods are shorter for living artists and postwar art.[13] This is in agreement with our earlier interpretation of more recent art as being more glamorous and with evidence that that glamour stocks (i.e., stocks with high market-to-book ratios) have higher turnover rates than low-priced value stocks (Hong and Stein 2007).

5.2. The Underperformance of Short-Term Transactions

We have shown so far that short-term transactions, or flips, are more common in booms and that price booms are followed by low returns. However, this does not necessarily imply that flips underperform on average. In fact, Lovo and Spaenjers (2018) argue the opposite; here, we replicate and extend their finding but show that it reverses after taking into account transaction costs.

We examine several specifications of the following basic equation:

$$\Delta p_{i,t}/h_{i,t} = \mu + \eta \log(h_{i,t}) + \varepsilon_{i,t}, \tag{11}$$

where $h_{i,t}$ is the holding period (in years) and the dependent variable $\Delta p_{i,t}/h_{i,t}$ is a measure of annualized returns. A positive η indicates that short-term transactions underperform longer-term transactions. Comparing transactions of different holding periods requires studying annualized returns.[14]

We are interested in how short-term transactions perform during booms and busts. If rational investors ride bubbles, we expect them to outperform the market during boom and bust cycles. To see if this is true, we decompose returns in Equation (11) into a market timing and an "art picking" component. The market timing component isolates the time-series effect and is informative about whether short-term investors buy art at the right time. The picking component captures the cross-sectional effect, that is, whether short-term investors buy art that appreciates the fastest, controlling for the time of purchase. To perform the decomposition, we first estimate a variant of Equation (2) whereby the left-hand side variable is the annualized round-trip return as in Equation (11). We then allocate the fitted values of this regression to each transaction. These fitted values capture the average return for all round-trip transactions on each purchase/resale year pair, that is, the market timing component. The residual in this regression captures the picking effect as it excludes the common component of time.

Of course, the timing of the decision to resell a work of art is an endogenous choice made by the investor, likely influenced by the price appreciation over the holding period. To address this concern, we also regress returns (as before, using (i) raw returns, (ii) market timing returns, and (iii) art picking returns) on the holding period of the *previous* owner—the lagged holding period, for short. If short-term transactions truly outperform, we should find a positive correlation between returns and the lagged holding period. To see this, consider an item i that appears twice at auction. It is first bought by collector A and then by collector B. Suppose A's transaction return is above average, and A's holding period is below average. It could be that A is a skilled flipper or that A decided to sell quickly after prices increased. We can separate the two effects by comparing B's return to A's holding period. If A's short holding period was simply a result of the appreciation of item i, then the item is neither underpriced nor overpriced relative to other artworks and A's holding period should be uncorrelated with B's returns. However, if A's return reflects skill, then the item is likely to be overpriced so that B would earn lower returns on average. Hence, working with lagged holding periods mitigates the concern that price increases would affect holding periods. Doing so, however, reverses the sign of our testable prediction; to make specifications comparable, we regress returns on the lagged holding period but with a minus sign so that a positive η indicates that short-term transactions underperform.

Whereas, elsewhere in the paper, we follow the literature by analyzing raw returns, an analysis on holding period returns requires taking into account transaction costs, which weigh heavily on returns over the short run but less so on returns over longer holding periods. We, thus, estimate a variant of Specification (11) assuming 20% transactions costs. As mentioned earlier, a 20% transaction cost is a lower bound as one also needs to account for the costs related to insurance, transportation, framing, and (country-specific) taxes.

Table 7 shows the partial correlations between annualized returns and holding periods. The first set of three columns presents the baseline results for returns before transaction costs and contemporaneous holding periods; the second set's results are based on lagged holding periods (with a minus sign); the last three columns introduce transaction costs in the baseline specification. In each specification, we present raw return estimates, which are dissected into return estimates based on market timing and art picking returns. The first three columns of Table 7 replicate and extend the results of Lovo and Spaenjers (2018) of a negative correlation between holding periods and annualized returns. Increasing the log holding period by one standard deviation (1.28) lowers future returns by 2.80% per year, which suggests that short-term traders

Table 7. Returns from Market Timing and Art Picking

	Baseline			Lagged holding period			After transaction costs		
Return	Raw	Timing	Picking	Raw	Timing	Picking	Raw	Timing	Picking
η	−2.18***	−1.00***	−1.18***	−2.49***	−0.16	−2.32***	6.50***	3.78***	2.72***
(Standard error)	(0.35)	(0.15)	(0.31)	(−0.62)	(−0.12)	(−0.61)	(0.37)	(0.16)	(0.34)
R^2	0.010	0.035	0.003	0.013	0.001	0.012	0.079	0.305	0.015
N	25,733	25,733	25,733	2,180	2,180	2,180	25,733	25,733	25,733

Notes. This table presents the estimates of the following regression:

$$\Delta p_{i,t}/h_{i,t} = \mu + \eta \log(h_{i,t}) + \varepsilon_{i,t},$$

where $\Delta p_{i,t}/h_{i,t}$ is the annualized return on transaction j over the holding period $h_{i,t}$ (in years). A positive η indicates that long-term transactions outperform short-term transactions. The table presents estimates for three specifications, each one including a decomposition of raw performance into market timing and art picking returns. The "baseline" specification assumes no transaction costs. The "lagged holding period" specification regresses returns of the current transaction on the holding period of the previous transaction (with a minus sign so that a positive η indicates that long-term transactions outperform). The "after transaction costs" specification assumes 20% transaction costs. Raw returns on the left-hand side are decomposed into a market timing and picking component by projecting raw returns into a combination of purchase and resale date fixed effects (see the text for more details). Standard errors are double-clustered at the year of purchase and the artist level.

***$p < 0.01$; **$p < 0.05$; *$p < 0.1$.

perform better, on average, before transaction costs. Market timing and art picking contribute equally to this performance. Although outperformance is potentially tainted by endogeneity, it is vindicated by our specification based on lagged holding periods. The correlations in the lagged holding period specification remain negative although the market timing component is now insignificant. Thus, an investor experiences lower returns, on average, when the investor purchases an item that was held over a short time period. The raw estimates have a similar magnitude as the baseline specification, alleviating reverse causality concerns. This is again in line with short-term traders being skilled. However, and crucially, the relationship between returns and holding periods reverses when taking into account transaction costs; short-term traders then do not outperform. Spreading out the transaction costs over long holding periods gives collectors an advantage. Both timing and picking components are positive and significant and are about thrice larger than the estimates without transaction costs.

In Online Appendix F, we also document that short-term transactions are significantly more volatile than longer transactions. For instance, the volatility of round-trip returns with holding periods less than one year is 60%, whereas the volatility of transactions between 9 and 10 years is 93%. We show that the difference is statistically significant with the variance ratio of round-trip returns slopping downward (see Online Table A.VIII).

To summarize, short-term returns are lower on average and more volatile. The average underperformance of short-term transactions is explained by transaction costs, which is consistent with prior evidence that stock investors who trade more have lower returns because they have to pay large fees for their trades (Barber and Odean 2000). Professionals in the art market often note that investors tend to overlook transaction costs, especially first buyers (see, e.g., Loader-Wilkinson 2010). Although we do not observe the identity of the individuals participating in the auction market, we suspect that such inexperienced investors drive a large portion of poor returns. This is in line with historical accounts of financial bubbles (e.g., Greenwood and Nagel 2009, Bayer et al. 2020). For instance, the early 2000s were characterized by the entry of many Russian investors on the art market, whereas as of 2005, Chinese and Indian buyers massively entered the art market (Kraeussl and Logher 2010, Renneboog and Spaenjers 2010). Besides, the volatility of short-term returns is so large that raw returns are often spectacular (as occasionally reported in the press), which leaves room to exceed transaction costs.

6. Relating the Facts to Theory

Throughout the paper, we document the following empirical regularities:

i. Lagged fundamentals predict returns (Section 3).

ii. Prices and volume are positively correlated (Section 4.1).

iii. Short-term transactions are procyclical (Section 4.1).

iv. Glamour assets are procyclical (Section 4.1).

v. Prices and volatility (price dispersion) are positively correlated (Section 4.2).

vi. In survey data, changes in disagreement are positively correlated with changes in volume and price dispersion (Section 4.3).

vii. In survey data, average beliefs and disagreement are positively related to lagged returns as well as lagged fundamentals (Section 4.3).

viii. During the late 1980s bubble, the correlation between returns and trading volume is stronger before the peak than afterward (Section 4.4).

ix. During the late 1980s bubble, volume peaked before prices (Section 4.4).

x. Short-term transactions underperform (Section 5).

We believe these findings can shed light on current theories aiming at explaining the joint behavior of asset prices and trading volume as well as speculative bubbles. Most of these theories have been developed to explain regularities in the stock and real estate markets but should also speak to the art market. In this section, we compare our empirical findings to the empirical predictions of a selection of papers. It is difficult to do justice to a huge literature, and we restrict our attention to papers that have explicit predictions about price and volume and/or belief dynamics.

6.1. Rational Learning

We begin the discussion with models that assume rational beliefs. Pastor and Veronesi (2009), in particular, propose a learning model in which bubble-like episodes arise in periods of high uncertainty. In their model, prices exhibit a boom–bust pattern when a new technology appears. At first, uncertainty is mostly idiosyncratic, pushing prices up. A bust occurs when uncertainty becomes more systematic, raising the risk premium for investing in the new technology.[15] The model produces a boom–bust pattern starting with good fundamentals as well as a positive correlation between prices and volume and between prices and volatility. In the context of the art market, it could indeed be that visual arts in general or some segments of the art market begin to appeal to new investors (e.g., new Russian oligarchs or the new Chinese wealthy) and that rational investors try to learn about whether new investors will indeed purchase art (or, in the Pastor–Veronesi context, whether the technology is adopted by a broader investment audience) or not. However, it is not clear whether this model could generate the patterns in trading volume we document, for example, the procyclicality of short-term transactions or our survey findings.

6.2. Loss Aversion

Loss aversion is a common explanation for the positive price–volume correlation in the real estate market (Genesove and Mayer 2001).[16] Barberis et al. (2001) propose a representative–agent model of the stock market with loss-averse investors. The model produces return predictability after a run-up in prices, after which agents become less risk averse. This is consistent with our predictability result (although we find that past fundamentals rather than past prices matter for art returns). However, the model is silent about trading volume because it features identical agents who have little reason to trade.

6.3. Heterogenous Beliefs/Utility

A plausible explanation for why people trade so much is that they hold different opinions. If short-sale constraints are also present, disagreement models predict that optimistic investors are willing to pay more than their own private value because they hope to sell to a greater fool before prices eventually fall (Harrison and Kreps 1978). Scheinkman and Xiong (2003) show how fluctuations in disagreement and short-sale constraints can generate comovement between prices, trading volume, and price volatility. Burnside et al. (2016) propose a model of speculative bubbles in which contagious social dynamics cause disagreement to fluctuate endogenously. Their framework generates a positive price–volume correlation but, by design, does not generate predictable returns.

Although in Scheinkman and Xiong (2003) and Burnside et al. (2016), agents have heterogeneous beliefs, in the context of art, agents do not need to disagree; their willingness to pay may simply differ. Lovo and Spaenjers (2018) propose a dynamic auction model in which rational agents' wealth and tastes are subject to random shocks that generate endogenous trading, and depending on their random preferences, agents emerge as collectors, investors, or speculators. Their model predicts that prices, volume, and the share of short-term transactions increase in good times when demand fundamentals are strong. However, because agents are strictly rational, there are no bubbles, and in particular, strong fundamentals do not predict the negative returns we document empirically. In addition, in their model, short-term traders act as rational arbitrageurs so that short-term trades earn higher returns on average. We find that this is only true when transaction costs are ignored.

6.4. Extrapolative Beliefs

So far, our discussion indicates that many of the stylized facts we document can be rationalized in a framework à la Scheinkman and Xiong (2003). However, although this framework is designed to speak of speculative bubbles, it is silent about why agents disagree. We argue earlier that our predictability results are consistent with models in which some investors display extrapolative beliefs. Extrapolation has long been proposed as an explanation for momentum/mean reversion patterns in asset prices (e.g., Cutler et al. 1991, Barberis et al. 1998).

We can distinguish models in which agents extrapolate past fundamentals from models in which they extrapolate past returns. We readily argue that our predictability evidence is more consistent with models in which agents over-extrapolate past demand fundamentals. For instance, Barberis et al. (1998) and Greenwood and Hanson (2015) propose that agents mistakenly believe that abnormally high fundamentals will persist into the future, leading to booms and busts in asset prices. However, these models feature identical agents, and the models do not speak of trading volume.

Barberis et al. (2018) and Defusco et al. (2020) explore the implications of return extrapolation and short-sale constraints in models of speculative bubbles. Barberis et al. (2018) propose that some investors extrapolate past price changes while also being attentive to on what other traders focus. As in Scheinkman and Xiong (2003), trading volume increases during a bubble because disagreement increases. However, disagreement increases endogenously because only a fraction of investors extrapolate, in line with the survey evidence in Section 4.3. Extrapolative and nonextrapolative investors trade against one another, which translates into the additional prediction of a strongly positive correlation between volume and past returns during bubble episodes, which we confirm using the late 1980s period in Section 4.4.

In Defusco et al. (2020), investors have homogenous beliefs but differ in their expected investment horizons. Price extrapolation arises endogenously in the model because agents neglect the sensitivity of market outcomes to others' beliefs as in Glaeser and Nathanson (2017). After a large demand shock, prices, volume, and the short-term buyer share initially all rise contemporaneously; in a second phase, prices continue to rise while volume and the short-term buyer share fall; finally, a bust occurs in which prices fall. In the model, the second quiet phase occurs because agents overestimate the level of the demand and expect it will continue to grow. This mistaken belief causes optimists to increase their reserve prices in spite of falling transaction volume. Thus, the model produces a positive price–volume correlation at relatively low frequencies and a lead–lag relation in which volume leads prices at higher frequencies. We empirically verify this prediction in Section 4.4.

6.5. Summary

Overall, our evidence depicts a coherent picture in which extrapolative investors are too optimistic in booms and too pessimistic in busts. We conclude that short-term speculators contribute to price bubbles as emphasized by the model of Defusco et al. (2020). We highlight the market frictions, short-sale constraints, and transaction costs that prevent arbitrageurs from correcting bubbles. In turn, short-sale constraints combined with differences in beliefs motivate speculative trading as predicted by Scheinkman and Xiong (2003) and Barberis (2018).

One stylized fact remains that we couldn't match to existing work on price and volume dynamics. We find that modern and contemporary art volume (i.e., glamour volume) is procyclical. In particular, it is not obvious why volume in this market segment rose so much in the late 1980s bubble. Historically, most bubbles have been backed by compelling and sensible stories. For example, the dot-com bubble fed on the idea that internet technology would bring dramatic productivity improvements; similar narratives were offered during the past railroad and electricity booms. We note in Section 4.4 that a popular story for the late 1980s bubble was that Japanese collectors were getting richer and increasingly eager to buy Impressionist art, which attracted speculators trying to anticipate the boom in Japanese demand. This story would have predicted that the boom concentrates on Impressionist art. Yet modern and contemporary art, not Impressionist art, experienced the largest boom in prices and trading volume. Arguably, this market segment includes artworks for which value is most uncertain. It is plausible that auction houses and speculators seized the opportunity of the Impressionist boom to feed the market with works of art that could be more difficult to sell in less euphoric times. This line of reasoning suggests a degree of substitutability between assets that may cause speculative dynamics to transit from one asset to another. We are not aware of a model that can produce such dynamics.

Another interesting aspect of our findings is that many ingredients of standard bubbles are absent from the art market. Historical accounts of financial bubbles emphasize the importance of radical innovation, new industries, and more generally the sentiment that "this time is different" (Kindleberger 1978, Greenwood and Nagel 2009, Pastor and Veronesi 2009). Another well-cited fact is the availability of cheap credit, leading to increased risk-taking and high leverage, which precipitates price busts through feedback loops and liquidity spirals (e.g., Brunnermeier 2009, Favilukis et al. 2016, Di Maggio 2017). None of these aspects seem to genuinely matter for art price formation. If anything, we show in Section 3.3 that credit booms predict negative art returns, and credit busts do not, which we interpreted as an additional element of evidence that biased beliefs are a central element of bubble formation.

7. Conclusion

The art market is subject to frequent booms and busts that seem difficult to reconcile with rational models in which people trade to consume. Textbook theory predicts that prices are forward looking, but we find that prices respond to current rather than future demand. We argue that price dynamics are consistent with art investors over-extrapolating demand growth into the future. Prices increase with current demand in a way that seems excessive as booms are followed by predictable busts. High prices are accompanied by high transaction volumes and a more than proportional increase in short-term transactions, which we interpret as trading frenzies given the substantial trading costs in the art market. During booms, trading concentrates on the works of postwar artists. Short-term transactions tend to underperform, but the volatility of short-term returns is much larger than transaction costs, which still leaves room for the spectacular profits occasionally reported in the press. In survey data, changes in disagreement are also positively correlated with changes in volume and price dispersion. Taken together, our findings are consistent with models in which extrapolative beliefs fuel speculation.

Acknowledgments
The authors are grateful to Utpal Bhattacharya, Patrick Bolton, Fabio Braggion, Brunella Bruno, Edouard Challe, Guillaume Chevillon, Alex Chinco, Pierre Collin-Dufresne, Gabriel Desgranges, Joost Driessen, Bernard Dumas, Rik Frehen, Jasmin Gider, Will Goetzmann, Kathy Graddy, Andreas Heinen, Harrison Hong, Frank de Jong, Pierken Jordain, David Laibson, Yuexin Li, Hao Liang, Roman Kräussl, Benjamin Mandel, Roberto Marfè, Kim Oosterlinck, Stavros Panageas, Lin Peng, Christophe Perignon, Laetitia Placido, Ludovic Phalippou, László Sándor, Olivier Scaillet, José Scheinkman, Kenneth Singleton, Christophe Spaenjers, Richard Taffler, Claudio Tebaldi, Patrick Verwijmeren, and Guillaume Vuillemey as well as participants of the 2012 French Finance Association Conference, the Eighth Art Market Symposium in Luxembourg, the Universite Libre de Bruxelles Art Markets Workshop in Brussels, the London Business School Summer Finance Symposium, European Finance Association 2016, Miami Behavioral Finance Conference 2016, ASU Sonoran Winter Finance Conference 2017, and Helsinki Finance Summit 2017 as well as seminars at THEMA, ESSEC, Tilburg University, Baruch College, Columbia University, Collegio Carlo Alberto, Ghent University, and Bocconi University for helpful comments and suggestions.

Endnotes

[1] Quoted in *The Wall Street Journal* (Peers 2008).

[2] Graddy and Hamilton (2021) study the effect of guarantees (both in-house and third party) and find that they have no significant impact on final prices. In the wake of the financial crisis, Sotheby's and Christie's substantially reduced providing in-house guarantees in late 2008. In this same year, Sotheby's lost about US$52 million in just one season when works did not sell at the price guaranteed to the seller, and it had to make up the difference (*Art and Design*, January 15, 2015). Since 2014, in-house price guarantees are cautiously used again.

[3] Information on the repeat sales sample is provided in Online Appendix A.1.3 (data appendix).

[4] We present results for the U.S. top wealth series to facilitate interpretation, but results are very similar when using the alternative series used in Table 1.

[5] To help alleviate concerns about inference in overlapping return regressions with persistent predictor variables (see, e.g., Ang and Bekaert 2007), we use Hodrick (1992) standard errors for the null of no predictability. In the repeat-sale specifications, we double-cluster standard errors based on the time of purchase and the name of the artist to account for residuals' cross-sectional correlation.

[6] We can also decompose further the variance of price changes at the transaction level. In Online Appendix E, we decompose this variance into systematic and idiosyncratic components.

[7] This result obtains under the simplifying assumptions that art price changes and the stochastic discount factor (SDF) are jointly normal and that the unobserved utility dividend is uncorrelated with the SDF.

[8] For consistency, we show results in which the dependent variable remains the log price change as in Table 2, but results are similar when we subtract the real risk-free rate from the left-hand side.

[9] Auction houses have incentives to increase volume when prices are higher, which could explain why prices and volume are correlated. However, the literature provides no evidence that auction houses engage in such behavior. Mei and Moses (2005) show that auction houses manipulate their presale estimates to increase their revenues but do not report evidence of time-variation in auction houses' behavior. Graddy and Hamilton (2021) find that more expensive items are more likely to have guarantees but do report that guarantees are more prevalent in booms. Ashenfelter and Graddy (2005) note that auction houses decided to increase fees in the aftermath of the 1980s bubble in a context of tepid prices and low volumes. Using data from Sotheby's annual reports, we find little evidence of a positive association between auction expenses (including marketing expenses) and trading volume.

[10] Tobias Meyer, who in 2006 was the director of Sotheby's contemporary art department worldwide, said to *The New York Times* (Vogel 2006), "Collectors want to beat the galleries at their own game … This insatiable need for stardom has made buying student work the art-world version of *American Idol*."

[11] We find consistent, albeit noisier, results when using data from the University of Michigan's survey of consumers about durable good purchases.

[12] To increase the sampling frequency over the bubble period, we present quarterly series, smoothed as four-quarter moving averages of the raw quarterly data. The exceptions are the Impressionist price and volume series, which are sampled annually because we work with a smaller sample of the works of Pierre-Auguste Renoir, Claude Monet, and Edgar Degas as did Hiraki et al. (2009).

[13] Holding periods are longer when we do not control for whether the artist passed away over the round trip. The premature death of an artist is associated with large returns (Pénasse et al. 2021), and anecdotal evidence suggests that investors speculate on death. But such speculation requires patience, hence, the longer holding periods for living artists.

[14] We find that annualized returns can be extreme, which is why we winsorize them at 5% and 95% to minimize the effect of potential outliers. Results based on unwinsorized returns are noisier but qualitatively similar.

[15] Frehen et al. (2013) present historical evidence consistent with that interpretation.

[16] An alternative explanation for this regularity is down-payment requirements in the mortgage markets (Stein 1995), but such requirements are essentially absent from the auction market.

References

Abreu D, Brunnermeier MK (2003) Bubbles and crashes. *Econometrica* 71(1):173–204.

Ang A, Bekaert G (2007) Stock return predictability: Is it there? *Rev. Financial Stud.* 20(3):651–707.

Ashenfelter O, Graddy K (2003) Auctions and the price of art. *J. Econom. Literature* 41(3):763–786.

Ashenfelter O, Graddy K (2005) Anatomy of the rise and fall of a price-fixing conspiracy: Auctions at Sotheby's and Christie's. *J. Competition Law Econom.* 1(1):3–20.

Bansal R, Yaron A (2004) Risks for the long run: A potential resolution of asset pricing puzzles. *J. Finance* 59(4):1481–1509.

Barber BM, Odean T (2000) Trading is hazardous to your wealth: The common stock investment performance of individual investors. *J. Finance* 55(2):773–806.

Barberis N (2018) Psychology-based models of asset prices and trading volume. Bernheim D, DellaVigna S, Laibson D, eds. *Handbook of Behavioral Economics: Applications and Foundations 1*, vol. 1 (North-Holland, Amsterdam, Netherlands), 79–175.

Barberis NC, Huang M, Santos T (2001) Prospect theory and asset prices. *Quart. J. Econom.* 116(1):1–53.

Barberis N, Shleifer A, Vishny R (1998) A model of investor sentiment. *J. Financial Econom.* 49(3):307–343.

Barberis N, Greenwood R, Jin L, Shleifer A (2018) Extrapolation and bubbles. *J. Financial Econom.* 129(2):203–227.

Baron M, Xiong W (2017) Credit expansion and neglected crash risk. *Quart. J. Econom.* 132(2):713–764.

Bayer P, Geissler C, Mangum K, Roberts J (2020) Speculators and middlemen: The strategy and performance of investors in the housing market. *Rev. Financial Stud*. 33(11):5212–5247.

Beggs A, Graddy K (2009) Anchoring effects: Evidence from art auctions. *Amer. Econom. Rev.* 99(3):1027–1039.

Brunnermeier MK (2009) Deciphering the liquidity and credit crunch 2007–2008. *J. Econom. Perspect.* 23(1):77–100.

Brunnermeier MK, Oehmke M (2012) Bubbles, financial crises, and systemic risk. Constantinides GM, Harris M, Stulz RM, eds. *Handbook of the Economics of Finance*, vol. 2 (Elsevier), 100–189.

Burnside C, Eichenbaum M, Rebelo S (2016) Understanding booms and busts in housing markets. *J. Political Econom.* 124(4): 1088–1147.

Campbell JY (1991) A variance decomposition for stock returns. *Econom. J. (London)* 101(405):157–179.

Campbell JY, Viceira LM (1999) Consumption and portfolio decisions when expected returns are time varying. *Quart. J. Econom.* 114(2):433–495.

Carlos A, Neal L, Wandschneider K (2006) Dissecting the anatomy of exchange alley: The dealings of stockjobbers during and after the South Sea bubble. Unpublished working paper, University of Illinois, Champaign.

Case KE, Shiller RJ (1987) Prices of single-family homes since 1970: New indexes for four cities. *New England Econom. Rev.* 45–56.

Chinco A, Mayer C (2016) Misinformed speculators and mispricing in the housing market. *Rev. Financial Stud.* 29(2):486–522.

Cochrane JH (2003) Stocks as money: Convenience yield and the tech-stock bubble. Hunter WC, Kaufman GG, Pomerleano M, eds. *Asset Price Bubbles: The Implications for Monetary, Regulatory, and International Policies* (MIT Press, Cambridge, MA), 122–167.

Cutler DM, Poterba JM, Summers LH (1991) Speculative dynamics. *Rev. Econom. Stud.* 58(3):529–546.

Davis J, Taylor AM (2019) The leverage factor: Credit cycles and asset returns. NBER Working Paper No. 26435.

Defusco AA, Nathanson CG, Zwick E (2020) *Speculative Dynamics of Prices and Volume* (Northwestern University and University of Chicago).

Di Maggio M (2017) Credit-induced boom and bust. *Rev. Financial. Stud*. 30(11):3711–3758.

Ellul A, Pagano M (2006) IPO underpricing and after-market liquidity. *Rev. Financial Stud.* 19(2):381–421.

Fama EF (2014) Two pillars of asset pricing. *Amer. Econom. Rev.* 104(6):1467–1485.

Favilukis J, Ludvigson SC, Van Nieuwerburgh S (2016) The macroeconomic effects of housing wealth, housing finance, and limited risk-sharing in general equilibrium. *J. Political Econom.* 125(1):140–223.

Frehen RGP, Goetzmann WN, Geert Rouwenhorst K (2013) New evidence on the first financial bubble. *J. Financial Econom.* 108(3):585–607.

Gabaix X (2012) Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance. *Quart. J. Econom.* 127(2):645–700.

Galbraith JK (1954) *The Great Crash 1929* (Houghton Mifflin Harcourt, Boston).

Galenson DW, Weinberg BA (2000) Age and the quality of work: The case of modern American painters. *J. Political Econom.* 108(4):761–777.

Genesove D, Mayer CJ (1997) Equity and time to sale in the real estate market. *Amer. Econom. Rev.* 87(3):255–269.

Genesove D, Mayer CJ (2001) Loss aversion and seller behavior: Evidence from the housing market. *Quart. J. Econom.* 116(4): 1233–1260.

Glaeser EL, Nathanson CG (2017) An extrapolative model of house price dynamics. *J. Financial Econom.* 126(1):147–170.

Goetzmann WN (1995) The informational efficiency of the art market. *Managerial Finance* 21(6):25–34.

Goetzmann WN, Renneboog L, Spaenjers C (2011) Art and money. *Amer. Econom. Rev.* 101(3):222–226.

Goyal A, Welch I (2008) A comprehensive look at the empirical performance of equity premium prediction. *Rev. Financial Stud.* 21(4):1455–1508.

Graddy K, Hamilton J (2021) Auction house guarantees for works of art. *J. Econom. Behav. Organ.* Forthcoming.

Graddy K, Loewenstein L, Mei J, Moses M, Pownall R (2014) Anchoring or loss aversion? Empirical evidence from art auctions. Discussion Paper No. 10048, Centre for Economic Policy Research.

Greenwood R, Hanson SG (2015) Waves in ship prices and investment. *Quart. J. Econom.* 130(1):55–109.

Greenwood R, Nagel S (2009) Inexperienced investors and bubbles. *J. Financial Econom.* 93(2):239–258.

Greenwood R, Shleifer A (2014) Expectations of returns and expected returns. *Rev. Financial Stud.* 27(3):714–746.

Greenwood R, Shleifer A, You Y (2019) Bubbles for Fama. *J. Financial Econom.* 131(1):20–43.

Harrison JM, Kreps DM (1978) Speculative investor behavior in a stock market with heterogeneous expectations. *Quart. J. Econom.* 92(2):323–336.

Hiraki T, Ito A, Spieth DA, Takezawa N (2009) How did Japanese investments influence international art prices? *J. Financial Quant. Anal.* 44(6):1489–1514.

Hodrick RJ (1992) Dividend yields and expected stock returns: Alternative procedures for inference and measurement. *Rev. Financial Stud.* 5(3):357–386.

Hong H, Stein JC (2007) Disagreement and the stock market. *J. Econom. Perspect.* 21(2):109–128.

Hong H, Kremer I, Kubik JD, Mei J, Moses M (2015) Ordering, revenue and anchoring in art auctions. *RAND J. Econom.* 46(1):186–216.

Jordà Ò, Schularick M, Taylor AM (2017) Macrofinancial history and the new business cycle facts. *NBER Macroeconomics Annual* 31:213–263.

Kahneman D, Tversky A (1974) Judgment under uncertainty: Heuristics and biases. *Sci.* 185(4157):1124–1131.

Kindleberger CP (1978) *Manias, Panics, and Crashes: A History of Financial Crises* (Macmillan, New York).

Kraeussl R, Logher R (2010) Emerging art markets. *Emerging Markets Rev.* 11(4):301–318.

Kräussl R, Lehnert T, Martelin N (2016) Is there a bubble in the art market? *J. Empirical Finance* 35:99–109.

Krigman L, Shaw WH, Womack KL (1999) The persistence of IPO mispricing and the predictive power of flipping. *J. Finance* 54(3):1015–1044.

Loader-Wilkinson T (2010) How art expenses stack up. *Wall Street Journal Europe*. (September 20), https://www.wsj.com/articles/SB10001424052748703959704575453162597091330.

Lovo S, Spaenjers C (2018) A model of trading in the art market. *Amer. Econom. Rev.* 108(3):744–774.

Mandel BR (2009) Art as an investment and conspicuous consumption good. *Amer. Econom. Rev.* 99(4):1653–1663.

Mei J, Moses M (2005) Vested interest and biased price estimates: Evidence from an auction market. *J. Finance* 60(5): 2409–2435.

Mei J, Scheinkman JA, Xiong W (2009) Speculative trading and stock prices: Evidence from Chinese AB share premia. *Ann. Econom. Finance* 10(2):225–255.

Mian AR, Sufi A (2009) The consequences of mortgage credit expansion: Evidence from the U.S. mortgage default crisis. *Quart. J. Econom.* 124(4):1449–1496.

Mian AR, Sufi A, Verner E (2017) Household debt and business cycles worldwide. *Quart. J. Econom.* 132(4):1755–1817.

Miller EM (1977) Risk, uncertainty, and divergence of opinion. *J. Finance* 32(4):1151–1168.

Ofek E, Richardson M (2003) DotCom mania: The rise and fall of internet stock prices. *J. Finance* 58(3):1113–1138.

Pastor L, Veronesi P (2009) Technological revolutions and stock prices. *Amer. Econom. Rev.* 99(4):1451–1483.

Peers A (2008) Will economy push art world off the ledge?—Auctions will test resilience of boom; prizing the prolific. *Wall Street Journal Europe* (May 8), https://professional.dowjones.com/factiva/.

Pénasse J, Renneboog L, Scheinkman J (2021) When a master dies: Speculation and asset Float. *Rev. Financial Stud*. Forthcoming.

Pénasse J, Renneboog L, Spaenjers C (2014) Sentiment and art prices. *Econom. Lett.* 122(3):432–434.

Pesando JE (1993) Art as an investment: The market for modern prints. *Amer. Econom. Rev.* 83(5):1075–1089.

Piazzesi M, Schneider M (2009) Momentum traders in the housing market: Survey evidence and a search model. *Amer. Econom. Rev.* 99(2):406–411.

Picinati di Torcello A, Petterson A (2014) Art & finance report 2014 (Deloitte Luxembourg and ArtTactic, Luxembourg).

Pownall RAJ (2017) *Tefaf Art Market Report 2017* (European Fine Art Foundation).

Renneboog L, Spaenjers C (2010) The iconic boom in modern Russian art. *J. Alternative Investments* 13(3):67–80.

Renneboog L, Spaenjers C (2013) Buying beauty: On prices and returns in the art market. *Management Sci.* 59(1):36–53.

Reyburn S (2014) For art collectors, the risk behind the high returns. *International New York Times* (June 9), https://www.nytimes.com/2014/06/09/arts/international/for-art-collectors-the-risk-behind-the-high-returns.html.

Scheinkman JA, Xiong W (2003) Overconfidence and speculative bubbles. *J. Political Econom.* 111(6):1183–1220.

Schularick M, Taylor AM (2012) Credit booms gone bust: Monetary policy, leverage cycles, and financial crises, 1870–2008. *Amer. Econom. Rev.* 102(2):1029–1061.

Schwert GW (2003) Anomalies and market efficiency. Constantinides G, Harris M, Stulz R, eds. *Handbook of the Economics of Finance* (Elsevier), 939–968.

Shiller RJ (1981) Do stock prices move too much to be justified by subsequent changes in dividends? *Amer. Econom. Rev.* 71(3): 421–498.

Shiller RJ (2000) *Irrational Exuberance* (Princeton University Press, Princeton, NJ).

Stein JC (1995) Prices and trading volume in the housing market: A model with down-payment effects. *Quart. J. Econom.* 110(2): 379–406.

Thompson D (2009) *The $12 Million Stuffed Shark: The Curious Economics of Contemporary Art* (Random House LLC).

Trucco T (1989) British pension fund sells $65.6 million in artworks. *The New York Times* (April 5), https://www.nytimes.com/1989/04/05/arts/british-pension-fund-sells-65.6-million-in-artworks.html.

Vogel C (2006) Warhols of tomorrow are dealers' quarry today. *The New York Times* (April 15), https://www.nytimes.com/2006/04/15/arts/design/warhols-of-tomorrow-are-dealers-quarry-today.html.

Wachter JA (2013) Can time-varying risk of rare disasters explain aggregate stock market volatility? *J. Finance* 68(3):987–1035.

Xiong W (2013) Bubbles, crises, and heterogeneous beliefs. Fouque JP, Langsam JA, eds. *Handbook for Systemic Risk* (Cambridge University Press), 663–713.

Xiong W, Yu J (2011) The Chinese warrants bubble. *Amer. Econom. Rev.* 101(6):2723–2753.