Asset Prices and Priceless Assets

Julien Penasse

Asset Prices and Priceless Assets

Proefschrift

ter verkrijging van de graad van doctor aan Tilburg University op gezag van de rector magnificus, prof.dr. Ph. Eijlander, en Université de Cergy-Pontoise op gezag van de president, prof. F. Germinet, in het openbaar te verdedigen ten overstaan van een door het college voor promoties aangewezen commissie in de Ruth First zaal van Tilburg University op dinsdag 2 december 2014 om 10.15 uur door

Julien Nicolas Guy-André Pénasse

geboren op 9 oktober 1982 te Eaubonne, Frankrijk.

PROMOTORES:

prof.dr. G. Desgranges

prof. dr. L.D.R. Renneboog

COPROMOTOR:

dr. E. Challe

OVERIGE LEDEN VAN DE PROMOTIECOMMISSIE:

dr. G. Chevillon

prof.dr. P. Collin-Dufresne

prof.dr. J.J.A.G. Driessen

dr. R.G.P. Frehen

prof.dr. F.C.J.M. de Jong

prof.dr. Scaillet

Acknowledgements

This dissertation concludes three years of hard work, dozens of pages of code, but above all a fantastic opportunity to meet and interact with bright and generous people, both at ESSEC, Université de Cergy-Pontoise, Tilburg, and more recently at Ecole Polytechnique.

I would like to start by thanking my advisors Luc Renneboog, Gabriel Desgranges and Edouard Challe. I have been privileged and honored to have them as advisors. I want to express my gratitude for their support, patience, and guidance. Five years ago, I was lucky to meet Edouard who encouraged me to pursue this Ph.D. I am deeply grateful to Edouard for giving me the idea to study volume in the art market: it goes beyond doubt that this thesis would have taken a different path without him. Although I decided to follow the empirical path of financial economics, both Gabriel and Edouard were always willing to give me feedbacks or brainstorm about the economics of art. I have been endlessly impressed by the precision and accuracy of Gabriel's remarks, and I must thank him for setting the bar high. I deeply enjoyed our long conversations at various coffee shops near Place d'Italie and was always inspired by our shared digressions on global games, Art, and sovereign debt crises. I am deeply indebted to Luc, who offered me to join Tilburg University. Luc had readily been a great co-author, and quickly became a fantastic advisor and a wonderful mentor. Since the very beginning — we had not even met — Luc has been of great advice, always eager to help me with professional, academic, or personal advice. I am immensely impressed by Luc's devotedness to his students, in spite of his many other responsibilities (including being interviewed on Dutch television!).

My deep gratitude also goes to the members of my dissertation committee. Hearing and reading their thorough feedback was a modesty lesson — a highly constructive one. I would like to thank Guillaume Chevillon, who followed me since the beginning at ESSEC. I thank Joost Driessen, Rik Frehen and Frank de Jong, for giving me extensive critical feedback on my ideas and pushing me to think harder about the contribution my research would make to the field. Finally, I thank Pierre Collin-Dufresne for his deep insights and for agreeing to travel to Tilburg to serve on my committee. I am also thankful to Olivier Scaillet for his valuable comments on this thesis.

The research presented in this thesis has been funded from several sources. I thank ESSEC Business School, the Université de Cergy-Pontoise, the LABEX ECODEC and Ecole Polytechnique. The last chapter of this thesis benefited from the support from the Research Initiative Long-term Asset Allocation at the Institut Louis Bachelier, sponsored by the Collège de France, the Caisse des Dépôts et Consignations and CNP Assurances. I would like to thank François Dezorme and Jean-Michel Beacco for inviting me to join this project, and Didier Janci, Isabelle Laudier and Stéphane Gallon for their insightful comments. I also thank Benjamin Poignard for his exceptional research assistance. Benjamin: I am looking forward to work with you again, now that you too embarked in the doctoral journey!

I also thank Roger Guesnerie for inviting me to collaborate with the Research Initiative EMMA of the Institut Louis Bachelier, sponsored by the Collège de France and the Crédit Agricole. I was lucky and honored to collaborate with him, Olivier Guéant, Jean-Michel Lasry and Marc-Antoine Autheman.

I am also thankful to Patrice Poncet, for encouraging and advising me in the very first days of this project. I would also like to thank Lorenzo Naranjo, Gorkem Celik, Patricia Charléty, André Fourçans, Estefania Santacreu-Vasut and Radu Vranceanu for their advice and their help before the job market. I am also grateful to Lina Prevost, who has provided outstanding support in all administrative and organizational tasks since the beginning of this Ph.D.

I would like to thank my co-author Christophe Spaenjers, for his valuable inputs. I have learned a great deal from him on how to write a paper, and the third Chapter of this thesis greatly benefited from Christophe's constructive skepticism. I look forward to working with him again in the future, and hope that he feels the same way.

This thesis was written in many different places. My gratitude goes to all my fellow PhD students in ESSEC and Cergy. Yes, I am talking about you Océane, Jo, Samia, Yana, Romain, Tania…I will always remember us talking about the weather (among other things). I wish you all an exciting and adventurous post-Ph.D. life, and I am looking forward for our paths to cross again. I will miss randomly visiting my office mates Samia (again!), Nourdine, Delphine, Esther, but also Raj, Max, Ran... Although I only spent a few weeks in Tilburg, I also want to thank Hao, Ran (again), Zorka, Marshall, Ferenc, Leila, and Larissa for making feel at home from the beginning.

I would like to express my gratitude towards my friends in Paris and elsewhere. Special thanks to the many non-economists who patiently listened to me struggling to explain in simple words what I was working on for all these years. Enfin, je voudrais remercier mes parents et frères pour leur soutien et leur patience, et pour m'avoir donné le goût du savoir. Mes derniers mots vont à la femme de ma vie, qui n'a jamais cessé de me soutenir dans ces moments de dur labeur avec une patience ineffable.[1] Elodie, si Tristan est à ton image, je serai le plus heureux du monde.

Julien Penasse

Vanves, September 2014

[1] "You say you're lookin' for someone
Who's never weak but always strong
To protect you an' defend you
Whether you are right or wrong"

Contents

Introduction

Understanding asset prices fluctuations is a central issue in financial economics. In a rational, no-bubble model, the present-value identity links the dividend-price ratio to future cash flows and future discount rates. In an efficient market, stock returns can be predictable if risk or risk premia vary over time. Stocks are "long duration" assets and small variations in discount rates can lead to large price fluctuations (see e.g. Cochrane, 2011). Thus stocks are not only predictable, but stocks prices move too much to be justified by subsequent dividends (Shiller, 1981).

There are many alternative or complementary views on return predictability. Alternative views generally emphasize the role of market imperfections — behavioral biases, market frictions and hence mispricing, in explaining return predictability. The title of Robert Shiller's Nobel lecture *Speculative Asset Prices* (Shiller, 2014), illustrates the central and controversial role of speculation and bubbles in price formation, two concepts that are explored in depth in the first three chapters of this thesis.

The controversy arises because there are too many models and not enough evidence to discriminate between these models. Schwert (2003) summarizes part of the profession's skepticism, doubting that behavioral models "have refutable predictions that differ from tests that have already been performed."

This doctoral thesis focuses on providing new evidence to help discriminating between theories. The first three chapters study various aspects of price formation in the art market. The final chapter proposes a Bayesian method to better extract information from the cross-section, and revisits international evidence on stock return predictability.

Chapter 1 exploits a unique survey data on the art community's confidence in the outlook of 21 "blue chip" artists. It shows that market sentiment, defined as the percentage of positive expectations minus the percentage of negative expectations, predicts short-term returns. It argues that fads may cause short-run price deviations from fundamental value, generating booms and busts in art prices. This paper is co-authored with Luc Renneboog and Christophe Spaenjers. It has been published in *Economics Letters* (Vol. 122, Issue 3, March 2014, pp. 432-434).

Chapter 2 implements a Mixed Data Sampling (MIDAS) modeling approach, which enables to predict year-end art returns using exogenous variables sampled at higher frequencies. The central ideal of this chapter is to measure how fast information diffuses in a decentralized market. Art is increasingly viewed as an investment vehicle, and it thus seems interesting to study the informational content of art prices. It takes about six months for art prices to incorporate information contained in the price of Sotheby's stocks. Trading volume and variables related to market sentiment have better short-term forecasting power. This paper has been presented at the sixth International Finance and Banking Society (IFABS) conference in Lisbon, the seventh Financial Risks International Forum in Paris and the third International Symposium in Computational Economics and Finance in Paris (ISCEF) in Paris.

Chapter 3 asks whether existing theories of return predictability can explain the large fluctuations in art prices, or whether these fluctuations reflect bursting bubbles. Speaking of bubbles is always challenging, even in the stock market. Quoting (Garber, 2001, p. 124)

> Before we relegate a speculative event to the fundamentally inexplicable or bubble category driven by crowd psychology, however, we should exhaust the reasonable economic explanations... "bubble" characterizations should be a last resort because they are non-explanations of events, merely a name that we attach to a financial phenomenon that we have not invested sufficiently in understanding.

Can we provide a reasonable explanation for art prices fluctuations? The starting point of this paper is that booms and busts in art prices generally come with parallel swings in trading volume. "High" prices are generally associated to large trading volume, a pattern which is reminiscent of many historical episodes of "speculative bubbles". Theories where agents are fully rational cannot provide a satisfactory explanation of the large trading volume that tend to accompany booms in asset prices: rational agents cannot agree to disagree. By contrast, behavioral biases such as overconfidence can generate disagreement among agents regarding asset fundamentals. When short-selling is costly, a buyer acquires the option to sell the asset to other agents when those agents have more optimistic beliefs, which can generate a significant bubble component in asset prices (Scheinkman and Xiong, 2003). Therefore, a high trading volume can signal the formation of a bubble and predict negative returns. This is exactly what we document in this paper. For example buying art in a "hot" market when volume is in the top first decile yields an average abnormal return of -3.5% per year. We also find that various measures of market sentiment and trading volume are significantly correlated with trading volume. This result is interesting because trading costs are huge in the art market, and also because traditional drivers of speculative bubbles (credit booms, leverage, . . .) are largely absent from the art market.

This paper is co-authored with Luc Renneboog. A previous version of this paper was awarded the French Finance Association Best PhD Workshop Presentation Award. It as been presented at the 29th Spring International Conference of the French Finance Association (AFFI) in Strasbourg, the 4th Annual Workshop on the History of Economics as Culture (Cergy-Pontoise) and in seminars at Université de Cergy-Pontoise, ESSEC and Tilburg University.

Chapter 4 addresses return predictability in the stock market. Most studies in the literature analyze comovement between stock returns and various "predictors," such as the dividend-price ratio. The slope coefficient of a regression of stock returns on lagged predictors is typically imprecisely estimated. This lack of precision limits the importance of return predictability, and often casts doubt on its significance. For example, international studies typically document a significant heterogeneity in the evidence across countries. This paper proposes to jointly consider the international evidence of return predictability. Intuitively, the premise is that the true parameters of international return processes share a common distribution. It is thus possible to learn about the common means and variances of the parameters and hence to formulate an informative prior that improves estimation precision. The model also makes a more efficient use of the cross-sectional correlation of the innovations. Once cross-sectional information is accounted for, the international evidence of return predictability appears much less fragmented than previously reported. In particular, there is reasonable evidence that the dividend-price ratio predicts both future returns and future dividend growth. Estimation risk is also substantially mitigated, so that stocks are typically safer in the long run than in the traditional framework.

A previous version of this paper, entitled “International Return Predictability and the Term Structure of Risk” has been presented at the 29th annual congress of the European Economic Association in Toulouse, the 31st Spring International Conference of the French Finance Association (AFFI) in Aix, the 23th Annual Meeting of the European Financial Management Association (EFMA) in Rome, the 2014 Annual Doctoral Conference of the Association for the Development of Research in Economics and Statistics (ADRES) in Paris, the Paris X Doctoral Conference in International Macroeconomics and Financial Econometrics in Nanterre, and in seminars at the EPFL@SFI, Cambridge University, Ecole Polytechnique, ESSEC, Université de Cergy-Pointoise and the Caisse des Dépôts et Consignations.

Bibliography

Cochrane, John H., 2011, Presidential address: discount rates, *Journal of Finance* 66, 1047–1108.

Garber, Peter M., 2001, *Famous first bubbles: The fundamentals of early manias* (MIT Press, Cambridge, MA).

Scheinkman, Jose A., and Wei Xiong, 2003, Overconfidence and speculative bubbles, *Journal of Political Economy* 111, 1183–1220.

Schwert, G. William, 2003, Anomalies and market efficiency, in George Constantinides, Milton Harris, and René Stulz, eds., *Handbook of the Economics of Finance*, north-holl edition, chapter 15, 937–972.

Shiller, Robert J., 1981, Do stock prices move too much to be justified by subsequent changes in dividends?, *American Economic Review* 71, 421–498.

Shiller, Robert J., 2014, Speculative asset prices, *American Economic Review* 104, 1486–1517.

Chapter 1

Sentiment and Art Prices[1]

Chapter Abstract

We hypothesize the existence of a slow-moving fad component in art prices. Using unique panel survey data on art market participants confidence levels in the outlook for a set of artists, we find that sentiment indeed predicts short-term returns.

I. Introduction

The art market shows remarkable boom-bust patterns. Returns to art investments are positively correlated in the short run (e.g., David et al., 2013), but may reverse in the longer run. Figure 1.1 illustrates the mean reversion in art prices around the 1990 art market peak, using data from Renneboog and Spaenjers (2013). It plots, for 13 art movements, the annualized real USD return between 1985 and 1990 against the horizontal axis. The corresponding returns between 1990 and 1995 are plotted against the vertical axis. A linear regression of the annualized returns between 1990 and 1995 on the returns between 1985 and 1990 results in a highly significant slope coefficient of 0.54 and an R-squared of 0.89.

[**Insert Figure 1.1 about here**]

The behavior of art prices is not well understood. We will argue that changes in art values cannot be fully accounted for by changes in fundamentals. Using unique new data, we will then examine whether variation in sentiment can help explaining art returns.

[1]The authors thank Anders Petterson of ArtTactic and Fabian Bocart of Tutela Capital for providing data. All errors are ours.

II. Fundamentals and fads

The fundamental value of a piece of art can be thought of as the sum of all discounted future ownership dividends (i.e., future flows of consumption services). In a representative-collector setting, this would imply that the correct price of artwork i at time t = 0 can be expressed as follows:

$$P_{i,0}^{Fund} = \sum_{t=1}^{\infty} \frac{E(D_{i,t})}{(1+r)^t} \tag{1.1}$$

The value of the future ownership dividends can be assumed to depend on the evolution of wealthreflecting the discretionary nature of luxury consumptionand on tastes. As tastes are slow-moving (Graddy, 2014), changes in (expected) financial wealth may be the prime determinant of changes in the fundamental value of art over the short run (Hiraki et al., 2009; Goetzmann et al., 2011). However, residuals from the regressions of art returns on economic fundamentals typically show the same bubble-like patterns as those reported for prices. Moreover, it is hard to see how wealth effects can explain the remarkable heterogeneity in returns across artists at any point in time.

The observation of booms and busts is consistent with the existence of fads. Camerer (1989) defines fads as mean-reverting deviations from intrinsic value caused by social or psychological forces. Market psychology is likely to affect beliefs about intrinsic value in the market for hard-to-value, impossible-to-short, and much-talked-about emotional assets such as art. Following Camerer (1989), we can formally incorporate a fad term F, capturing beliefs about the consumption services that will flow from the ownership of a piece of art, by adapting Equation (1.1) as follows:

$$P_{i,0}^{Fad} = F_{i,0} \sum_{t=1}^{\infty} \frac{E(D_{i,t})}{(1+r)^t} \tag{1.2}$$

with F having a mean of one and changing over time as follows:

$$F_{i,t+1} = C_{t+1} \times F_{i,t} + \epsilon_{i,t+1} \tag{1.3}$$

where C is a parameter that determines whether the fad is growing ($C > 1$) or decaying ($C < 1$), and ϵ is a zero-mean and independent error term. According to Equation (1.2), faddish beliefs should be positively related to price levels. Furthermore, if we assume that the fad is indeed mean-reverting (and not a rational growing bubble, for example), the magnitude of the fad component should be negatively related to longer-term returns.

It is of course impossible to directly observe the expected dividends from art ownership and therefore whether a fad component exists for any individual artwork or artist. Yet, it is clear that a growing fad component ($C > 1$) should translate in a subjective expectation of observing much higher prices in the near future. We will call such expectations of higher prices high sentiment from now on. We expect high sentiment to be accompanied (and immediately followed) by increases in price levels. An extended period of high sentiment signals that a fad has been growing for a long time, and should be related to relatively low returns over the long run. By contrast, a decaying fad component ($C < 1$) implies subjective expectations of price depreciation, i.e., low sentiment. We expect low sentiment to be accompanied by decreases in prices in the short run. Extended periods of low sentiment should predict relatively high financial returns over the long run.

Renneboog and Spaenjers (2013) construct a market-wide proxy for sentiment, and find a relation between sentiment and next-year returns. However, their measure can only exploit time-series variation in beliefs. By contrast, in this paper, we use a unique panel data set containing information on sentiment at the level of the individual artist.

III. Data

ArtTactic, a London-based art market research firm, has surveyed a pool of art market playerscollectors, auction houses, dealers, etc.on their short-term confidence in a set of artists on a semi-annual basis since November 2005. The question asked by the firm is the following one: How do you feel about the artists market in the next 6 months? Possible answers are positive, neutral, and negative. We have data on the variation in the art market communitys confidence in 70 American and European post-war and contemporary artists. ArtTactic started with a list of 24 contemporary artists (e.g., Damien Hirst, Richard Prince) in 2005; 16 other artists were added later. The company has also surveyed art market professionals confidence in 30 blue chip post-war artists (e.g., Andy Warhol, Francis Bacon) since early 2008. The latest data used in this paper stem from November

2012. For each artist i and each period t, we compute a sentiment measure by subtracting the percentage of negative answers from the percentage of positive responses:

$$Sentiment_{i,t} = (\%Positive - \%Negative)_{i,t} \tag{1.4}$$

We find substantial cross-sectional variation in our sentiment measure. A linear regression of sentiment on semester dummies results in an R-squared of not more than 0.17. By contrast, sentiment is persistent: a regression of our sentiment variable on artist fixed effects yields an R-squared of 0.52, and the autocorrelation coefficient equals 0.77.

Figure 1.2 shows the evolution of the average level of sentiment per half-year since the second half of 2005. The most striking aspect of Figure 1.2 is probably the sharp drop in sentiment over the second half of 2008. The survey of November 2008 was the only one for which the proportion of negative outlooks exceeded the proportion of positive outlooks on average.

[**Insert Figure 1.2 about here**]

We merge our sentiment data with semi-annual artist-specific price indexes for the period 2004-2012 from Tutela Capital, a provider of art market information. We drop all artists with less than 20 sales during the first half of 2004 from the sample, because estimates of price indexes are typically noisy when based on few data. This exclusion restriction leaves us with 21 artistsnot a large sample, but still an improvement over data sets that only include time-series information.

IV. Results

We examine the predictive power of sentiment for short-term returns. In the first column of Table I, we regress the log price change for artist i between semester $t-1$ and semester t on the sentiment level for artist i near the end of semester $t-1$. In the second column, we control for the returns on the S&P 500 over the six-month periods leading up to the ends of periods $t-2$, $t-1$, and t, as changes in financial wealth may affect the fundamental values of artworks. In the third column, we add period fixed effects, to absorb changes in economic fundamentals over time, which should affect all artists similarly. In other words, we examine whether, cross-sectionally, higher-than-average sentiment is related to higher-than-average returns. In each case, we cluster standard errors both by artist and by time period.

[**Insert Table I about here**]

The regression results in the first three columns of Table I show that higher sentiment levels are indeed correlated with faster price appreciations. This result holds when controlling for the returns on equities, and when including period fixed effects in our model. Moreover, the results are also economically significant. For example, the coefficient of 0.11 found in the second and third regression model implies that an increase in the level of sentiment of 0.29 (the standard deviation of our sentiment variable across the full set of artists and time periods) is associated with an increase in the half-yearly log return of more than 3 percentage points. To mitigate concerns that our results are driven by reverse causalityfor example, price trends starting in semester $t-1$ could affect sentiment near the end of $t-1$ — we also repeat our regression models using sentiment in period $t-2$. The results are reported in the next three columns of Table I, and are very similar to those reported before.

V. Conclusion and discussion

Using unique survey data on the art communitys confidence in the outlook for a set of artists, we find substantial evidence that high sentiment positively predicts art returns over the short run, in line with our expectations. Unfortunately, our panel data set currently does not allow a robust analysis of whether extended periods of high sentiment predict low long-term returns. Also a study of the factors that drive fads in the art market is left for future research.

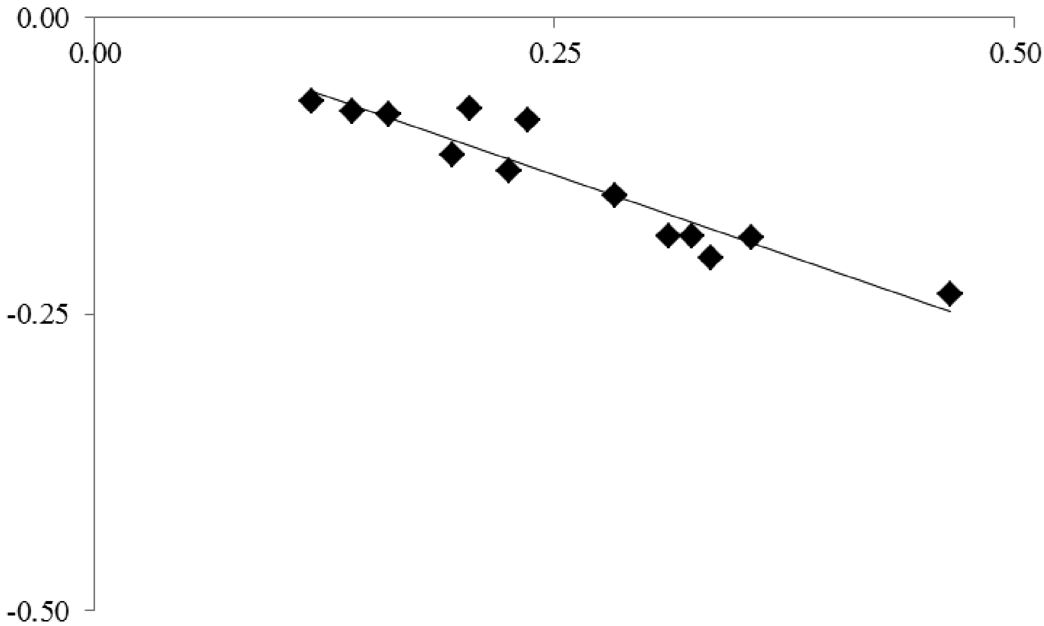

Figure 1.1: Figure 1 illustrates the mean reversion in prices around the art market peak in 1990. It plots, for 13 art movements, the annualized real USD return between 1985 and 1990 against the horizontal axis, and the corresponding returns between 1990 and 1995 against the vertical axis. Also the trendline is shown.

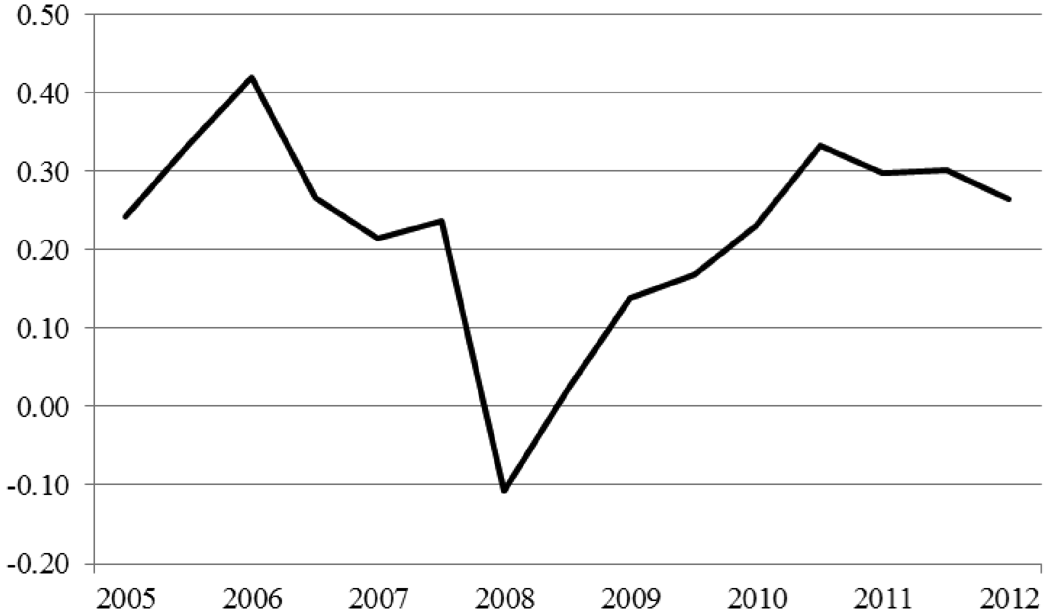

Figure 1.2: Figure 1.2 shows the evolution of average sentiment between the second half of 2005 and the second half of 2012 for the artists considered in this study. Sentiment is measured using ArtTactic survey data on the short-term confidence in the market for each artist. A positive value signals a higher proportion of positive than negative views.

Table I: Table I shows the results of regressions that relate changes in the artist log price indexes to lagged sentiment levels and controls. Price indexes are provided by Tutela Capital. Sentiment is measured using ArtTactic survey data on the short-term confidence in the market for each artist.

	Δ Log price index					
	(1)	(2)	(3)	(4)	(5)	(6)
Sentiment_{-1}	0.15 **	0.11 **	0.11 ***			
	(0.06)	(0.04)	(0.03)			
Sentiment_{-2}				0.10 **	0.12 ***	0.12 ***
				(0.05)	(0.03)	(0.04)
(Δ Log equities)		0.37 ***			0.42 ***	
		(0.02)			(0.06)	
$(\Delta \text{ Log equities})_{-1}$		0.24 **			0.32 ***	
		(0.12)			(0.12)	
$(\Delta \text{ Log equities})_{-2}$		0.12			0.14 *	
		(0.07)			(0.08)	
Period fixed effects?	No	No	Yes	No	No	Yes
N	211	211	211	190	190	190
R^2	0.04	0.12	0.18	0.02	0.11	0.16

Bibliography

Camerer, Colin, 1989, Bubbles and fads in asset prices, *Journal of Economic Surveys* 3.

David, Géraldine, Kim Oosterlinck, and Ariane Szafarz, 2013, Art Market Inefficiency, *Working Papers CEB* 32.

Goetzmann, William N., Luc Renneboog, and Christophe Spaenjers, 2011, Art and money, *American Economic Review* 101, 222–226.

Graddy, Kathryn, 2014, Taste Endures! The rankings of Roger de Piles (1709) and three centuries of art prices, *Journal of Economic History* 73, 766–791.

Hiraki, Takato, Akitoshi Ito, Darius A. Spieth, and Naoya Takezawa, 2009, How did Japanese investments influence international art prices?, *Journal of Financial and Quantitative Analysis* 44, 1489.

Renneboog, Luc, and Christophe Spaenjers, 2013, Buying beauty: on prices and returns in the art market, *Management Science* 59, 36–53.

Chapter 2

Real-time Forecasts of Auction Prices[1]

Chapter Abstract

This paper uses the Mixed Data Sampling (MIDAS) modeling approach to forecast aggregate prices in the fine arts auction market. Art price indices are released to the public on a low frequency basis, and MIDAS regressions allow to forecast year-end returns using higher frequency variables: macro-financial variables and proxies for market sentiment. It takes about six months for art prices to incorporate information contained in the price of Sotheby's stocks. Variables related to art market sentiment have better explanatory power in the short-term. Out of sample, art prices changes are largely unpredictable, in line with similar evidence in the stock market. These findings suggest that macro-financial information diffuses only gradually into the art market.

I. Introduction

Asset prices should reflect currently available information; markets that fail to aggregate current information are generally considered to be inefficient. Some markets differ from this theoretical paradigm in that traded goods can be heterogeneous and be traded infrequently, which greatly complicates information aggregation in a timely manner. The auction market for fine arts is one such market. Art price indices exist, but are released to the public on a low frequency basis, typically once a year. Market information is thus disseminated to collectors and amateurs in the form of auction outcomes or interim price

[1]I would like to thank Edouard Challe, Guillaume Chevillon, Pierre Collin-Dufresne, Joost Driessen, Frank de Jong, Luc Renneboog and Olivier Scaillet for their helpful comments and suggestions. I also thank the conference participants to the sixth IFABS (Lisbon), third ISCEF (Paris) meetings and the 7th Paris Financial Risks International Forum. Last, I am grateful to Eric Ghysels for providing Matlab routines for fitting MIDAS models. All remaining errors are my own.

indices. Such information is generally fee-based, so that the true level of prices is generally not common knowledge.

This paper considers the problem of information aggregation from an empirical point of view. In the absence of a reliable price index, markets participants may try to proxy for potentially unobserved information. For example, the hedge fund and art collector manager Jim Chanos recently claimed that Sotheby's stock is a good proxy for art prices, and that one could use it to hedge his art collection.[2] If an arbitrage relation between Sotheby's stock and auction prices genuinely exists, market participants should be able to quickly extract current information from Sotheby's stock prices.

Assessing the ability of a "high" frequency variable to proxy for a low frequency price index requires a model where data is sampled at different frequencies. The Mixed Data Sampling (MIDAS) approach addresses the difference in sampling frequencies between variables by employing a weighted time aggregation. The weights are chosen to be functions of the elapsed time and an estimated vector of hyperparameters. This framework allows me to perform regressions with leading variables, and therefore to forecast price changes on short horizons. Coming back to our forecasting example, the MIDAS approach will construct a weighted average of recent Sotheby's stock returns. This composite variable will be used as a leading indicator of year-end art prices.

I will focus on year-end prices and vary the forecasting horizon. This will allow me to measure the quality of candidate proxies for art prices. A "good" proxy, i.e. a variable that is significantly correlated to art prices, should have an explanatory power than decreases with forecast horizon. I show that if a variable leads art prices with a significant lag, however, its explanatory power may *increase* with horizon. In addition to the practical objective of forecasting, this methodology will therefore measure the informational content of art prices.

Beyond Sotheby's stock, where could one find information about art prices? Unlike traditional financial assets such as stocks, the prices of art objects are not "anchored" by the flow of dividends that are expected to accrue to the shareholder. A growing literature studies art and collectibles as assets.[3] The fundamental value of a work of art can be seen

[2]Quoted on CNBC (Frank, 2014).

[3]See Burton and Jacobsen (1999), Ashenfelter and Graddy (2003) and Goetzmann et al. (2014) for reviews of the literature.

as the expected discounted value of future "utility" dividends, which is the rent one would be willing to pay to own this work of art over a given time frame. The previous literature therefore identifies two major factors that affect art prices. Most papers emphasize the importance of macroeconomic fundamentals. Although tastes about individual artists may fluctuate randomly, aggregate art prices are largely driven by demand from the wealthy. This hypothesis has been empirically supported by, e.g., Hiraki et al. (2009), Goetzmann et al. (2011), Pownall et al. (2013) and theoretically by Mandel (2009). Several recent studies, on the other hand, suggest that prices can deviate substantially from fundamental value. Measures of market sentiment have been shown to predict art returns (Renneboog and Spaenjers, 2013; Penasse et al., 2014). Penasse and Renneboog (2014) show that the art market tends to be characterized by episodes of trading frenzies that coincide with booms and busts in prices.

This paper makes use of an extensive data set containing information on 141,638 sales of works of art by 87 major artists. The sample spans 1954 to 2010, but I focus on the 1973-2010 period, in order to replicate the updating of a hypothetical annual price index in real time. In addition to Sotheby's stock, I consider as potential predictors macro-financial variables and variables related to market sentiment: volume and the "anxious index" produced by the Survey of Professional Forecasters. The out-of-sample forecasting ability of these variables is compared to two benchmarks: a "tracking" index that makes use of within-year prices and a "historical average" benchmark, that ignores within-year information and forecasts future prices to grow at a constant rate.

A unique feature of the auction market is the presence of presale estimates. Such estimates are provided by auction house experts, who by their very position are the most likely to provide accurate forecasts of current art prices, and therefore aggregate information. This paper shows that auction house experts fail to provide unbiased estimates of current price levels. Estimation errors are persistent and tend to comove with prices, suggesting experts systematically underestimate price volatility.

Turning to the main analysis, an interesting pattern emerges. In the short-run, the real-time informational content of financial variables is relatively poor. In particular, Sotheby's stock is almost uncorrelated to art returns up to a six-month horizon, but tends to gain predictive power as horizon increases. Sentiment-related variables, on the other hand, have better short-term forecasting power. Assessing out-of-sample forecast performances, however, I find art returns to be largely unpredictable, even in the short term. With the exception of Volume and for horizons of up to six months, MIDAS predictions fail to consistently beat the historical average. In spite of a large documented correlation with art market returns, equity returns provide the worst predictions for all horizons except 9 months.

This finding is consistent with gradual information diffusion. If art prices react to macroeconomic fundamentals with a substantial lag, then the explanatory power of these variables should increase with forecast horizon. This pattern of predictability is likely to be due to market segmentation. If stockholders only infrequently participate to the auction market and if art market participants are inattentive to macroeconomic information, the latter will only diffuse slowly into art prices. The impossibility to observe aggregate art prices in real time is also likely to exacerbates the inability of the auction market to quickly react to macroeconomic information. In contrast, variables related to market sentiment are more likely to reflect art market participants' expectation, and thus to have better short-term forecasting power.

A substantial literature studies theoretically and empirically how information diffuses into asset prices. Hong and Stein (1999) posit that gradual diffusion of information among investors can explain momentum and return predictability. Hong et al. (2000) provide empirical support for the gradual-information-diffusion model of Hong and Stein, and find that information, and especially negative information, diffuses slowly into prices. Pollet and DellaVigna (2007) find that demographic information can be used to predict future stock returns and interpret their results in terms of information diffusion. Several studies document lead-lag relation between prices across stocks (e.g. Lo and MacKinlay 1990; Brennan et al. 1993), industries (Hou, 2007; Menzly and Ozbas, 2010) and countries (Rapach et al., 2013). Huberman and Regev (2001) show a firm's stock price soared on the release of information in the *New York Times* that had been published in *Nature* and various popular newspapers 5 months earlier. Cohen and Frazzini (2008) present evidence that stock prices adjust with a lag to news involving economically related firms, inducing predictable returns. This paper contributes to this literature by showing how mixed frequency regressions can be used to test information diffusion, and is the first to do so in an illiquid market.

This article is also related to a strand of studies investigating market efficiency in the fine arts auction market. Frey and Eichenberger (1995) argue that the art market is prone to behavioral anomalies because many collectors are not profit oriented. Penasse et al. (2014) and Penasse and Renneboog (2014) stress the role of market sentiment in price formation, in particular during boom-bust episodes. Strict tests of market efficiency have provided conflicting evidence. Erdos and Ormos (2010) find that weak-form efficiency cannot be rejected (at least for the past 64 years), while David et al. (2013) reject the same hypotheses, pointing that reserve prices give informational superiority to market insiders.

The remainder of the study is organized as follows. Section II presents a simple model that describes the empirical design and generates the main testable predictions. Section III describes the art auction data and the construction of variables. Section IV provides evidence of systematic error in expert presale estimates. The Mixed Data Sampling methodology is introduced in Section V. Section VI and VII present the main estimations and discusses out-of-sample ability of the proxies. Section VIII concludes.

II. Model

My model consists in a four-date economy, $t = 0, 1, 2, 3$. Denote R_t the cumulated price changes, or returns, at time t. Investors may not observe R_t.[4] Rather, they trade heterogeneous goods, which individual value is strongly correlated to R_t. In my empirical setup, R_t corresponds to cumulated changes in a price index, i.e. a weighted average of individual trade prices, after individual quality has been controlled for. I assume that $R_t = R_{t-1} + u_t$, where $R_0 = 0$ and u_t is i.i.d. $\mathcal{N}(0, \sigma_R^2)$.

We are interested in forecasting R_3 at times $t = 1, 2$. I will denote this exercise as *nowcasting*, because some information is readily revealed about R_3 when t is greater than 0. If one could observe the true value of R_1 (R_2), he could use it to forecast R_3. It is easy to verify that the R-squared of such regression is 1/3 (respectively 2/3).

Investors that do not observe aggregate returns may try to nowcast them using alter-

[4] I provide evidence in Section IV that even auction house experts may never perfectly observe current prices.

native variables. I consider two variables. The first, V_t, is contemporaneously correlated to R_t while the second, S_t, has forecasting power over R_t:

$$\begin{aligned} V_t &= \gamma R_t + v_t \\ S_t &= \theta R_{t+2} + w_t \end{aligned}$$

where v_t is i.i.d. $\mathcal{N}(0, \sigma_V^2)$ and w_t is i.i.d. $\mathcal{N}(0, \sigma_S^2)$. V_t can correspond to trading volume. Penasse and Renneboog (2014) shows that price and volume are strongly correlated in the art market and argues that volume can be seen as a measure of market sentiment (see also the discussion in Section III.C). S_t can be seen as Sotheby's stock price. In a frictionless and perfect-information economy, S_t and R_t would be contemporaneously correlated, because Sotheby's dividend is a function of future art prices (multiplied by volume), and because both S_t and R_t ultimately depend on the future demand for art. Rather, I assume that art prices react with a substantial lag to the information contained in Sotheby's stock price. This can occur if investors participate in either the stock or art market, and if art investors have limited ability to process information from the stock market (see e.g. Hong et al. 2007).

Testable prediction. *The forecasting power of V_t (respectively S_t) decrease (increase) with forecasting horizon.*

The R-squared for forecasts using either V_t or S_t are $R^2_{V_1} = \frac{\gamma^2\sigma_R^2}{3(\gamma^2\sigma_R^2+\sigma_V^2)}$, $R^2_{V_2} = \frac{4\gamma^2\sigma_R^2}{3(2\gamma^2\sigma_R^2+\sigma_V^2)}$, $R^2_{S_1} = \frac{3\theta^2\sigma_R^2}{3\theta^2\sigma_R^2+\sigma_S^2}$ and $R^2_{S_2} = \frac{3\theta^2\sigma_R^2}{4\theta^2\sigma_R^2+\sigma_S^2}$. It is easy to verify that $R^2_{V_1} < R^2_{V_2}$ and $R^2_{S_1} > R^2_{S_2}$. In this framework, it is necessary for S_t to lead R_t by two periods for the forecasting power to increase with horizon. For example, if instead volume leads R_t by one period

$$\mathcal{V}_t = \gamma R_{t+1} + v_t,$$

then $R^2_{\mathcal{V}_2} = \frac{3\gamma^2\sigma_R^2}{3\gamma^2\sigma_R^2+\sigma_{\mathcal{V}}^2}$ is always larger than $R^2_{\mathcal{V}_1} = \frac{4\gamma^2\sigma_R^2}{3(2\gamma^2\sigma_R^2+\sigma_V^2)}$.

III. Data

A. Data set

The data set used consists of 141,638 sales concerning 87 major artists from 1954 to 2010, obtained from the online database Art Sales Index. The artists are mainly from the nineteenth and twentieth century. The data set mostly consists of works on paper (42%), paintings (27%) and reproducible items such as etchings or lithographs (26%). Prices are hammer prices, exclusive of transaction costs. For each item, information has been collected on the work itself and the location of the sale. The following information is included: the price of the painting, the name of the artist, its title, place and date of sale, dimensions, technique and medium used, signature, and date. Some items are numbered, signaling there are multiple version of the same work, e.g. different impressions of the same lithograph. Finally, pre-sale estimates are available from 1995, for 72,429 sales.

B. Art price indices

Hedonic regression

I construct art price indices by applying hedonic regressions on art prices. Hedonic regressions are a popular methodology for constructing constant-quality price indices for infrequently traded goods such as houses or collectibles. Hedonic models seek to keep constant the objective characteristics of each work of art by including a small number of hedonic characteristics (name of the artist, medium, etc.). It can be written as:[5]

$$\ln(P_{kt}) = \sum_{t=1}^{T} p_t \delta_t + \sum_{i=1}^{K} \alpha_i x_{i,kt} + \epsilon_{kt} \tag{2.1}$$

where $\ln(P_{kt})$ represents the natural logs of prices and δ_t is a time dummy that takes a value of 1 for an artwork sold at time t. The coefficients $p_{t=1,\ldots,T}$ lead to an index of (log) prices. The vector $x_{i,kt}$ is composed of variables coding the characteristic of each object

[5] Art price indices can alternatively be constructed through repeat-sales regression (RSR), using only purchase and sale prices of individual artworks traded at two distinct moments in time. The inconvenient of such a procedure is that it drastically reduces the number of observations, so that the estimate of the index can be more volatile and prone to autocorrelation. For a comprehensive discussion of the pros and the cons of both types of indices see, e.g. Ginsburgh et al. (2006).

k at time t: the name of the artist, the technique used, the size of the artwork and other variables affecting the value of an artwork (is it signed? dated? is it part of a series?).

Annual price index

Table I presents the results of a hedonic regression on the full sample. In order to replicate the updating of an actual art index in real time, I estimate Equation (2.1) dynamically. The first value of the index, denoted $p_{1973|1954:1973}$ is obtained from a hedonic regression based on data available in 1973 (20 years of data are available at that point). The following values are constructed using an expanding window (i.e. 1954-1974, 1954-1975, etc.). For the purpose of this paper, I will treat this annual index as the "true" level of the art market. I assume that it is released on the last day of each year. Finally, the series is converted to returns by taking the log-difference of the price series: $r_t = p_{t|1954:t} - p_{t-1|1954:t-1}$.

[Insert Table I about here]

Tracking index

The hedonic method is also used to produce an index tracking updates during the year. This tracking index will serve as benchmark for out-of-sample forecasts, and corresponds to the interim values R_1 and R_2 in the notation terms of Section II. It is constructed as the annual index, except that it only uses sales that occurred up to the update date. I update this index every quarter: the first quarter value is constructed using all sales between the beginning of January to the end of March and is released on March 31, the second quarter uses sales from January to June and is released on June 30, etc. The fourth quarter of the tracking index therefore coincides with the annual index. The annual and tracking indices are reproduced in Figure 2.1. The art market is known to be seasonal since, for calendar reasons, major sales occur before and after summer. The tracking index indeed seems to present a form of noise or seasonality for some years, although not systematically. Such a time-varying seasonality occurs because artwork heterogeneity cannot be perfectly observed and there is no simple way to address it.[6] This noise in the

[6]Including month dummies only makes things worse by increasing the number of regressors.

tracking index reflects the intrinsic limits of "high" frequency indices, thus motivating the use of proxies that I introduce in the next section.

[Insert Figure 2.1 about here]

The tracking index is then used to construct forecasts of annual returns r_t. I build these forecasts under the assumption that art returns are unpredictable. This assumption implies that the best forecast of returns between the time of the forecast and the end of the year is the average of past returns. I compute this average as

$$\bar{r}_t = \frac{1}{t} \sum_{\tau=1}^{t} r_\tau. \tag{2.2}$$

so that the within-year forecasts of annual returns write[7]

$$\hat{r}_{t|t-q/4,p} = p_{t-q/4} - p_{t-1} + (1 - q/4)\bar{r}_{t-1}. \tag{2.3}$$

Observe that within-year observations are denoted by a fractional index: $\hat{r}_{t|t-q/4,p}$ is the forecast of year-end return q quarters ahead and $p_{t-q/4}$ is the value of the tracking index lagged by q quarters. Finally, p_{t-1} is the value of the true annual art index at the end of the previous year. Equation (2.3) says that the year-end forecast is equal to current return and a return equal to historical average from the forecast date. As the year progresses, the forecasting period becomes shorter, and hence forecasts become more accurate. I will compare these forecasts to predictions from the MIDAS framework.

C. Predictors

I consider three monthly financial variables as potential predictors of art returns. Sotheby's stock price is the natural starting point of our analysis and is downloaded from Bloomberg.[8] The Sotheby's series starts in May 1988, while the remaining financial variables span 1973-2010. The latter are macro-financial variables motivated by the art market and asset return predictability literatures. Stock market returns have been shown to lead the art market, typically by one year (Goetzmann, 1993; Chanel, 1995). A popular

[7] I drop the subscript 1959:t for notational simplicity.

[8] Sotheby's main competitors, Christie's and Phillips are not listed at the time of writing this paper.

interpretation is that equities drive the art market because the demand for art increases with the wealth of collectors. I obtain equity returns from Morgan Stanley Capital International database (MSCI World). Finally, bond returns reflect risk premia as well as inflation and growth expectations, and may therefore contain information about future art prices. I use the Datastream 10-year US government benchmark to construct bond returns, using the duration loglinear approximation described in Chapter 10 of Campbell et al. (1997).

I also consider three variables related to market sentiment. It is common to consider trading volume – or liquidity – as a good proxy for market sentiment (Baker and Stein, 2004). Trading volume has also been shown to covary with art returns (Korteweg et al., 2013; Penasse and Renneboog, 2014). Theoretical explanations of correlation between returns and volumes include disagreement among speculators (Miller, 1977; Scheinkman and Xiong, 2003) or the presence of irrational investors (Baker and Stein, 2004). Since short selling is impossible in the art market, the holder of a given artwork will generally be the most optimistic about its value. The opinions of the pessimists will thus fail to be incorporated into prices, which will then only reflect the opinions of the optimists. As a result, an investor can be willing to pay more than his own private value for a painting because he expects that, in the future, there may be other investors that value the painting more than he does. The difference between his willingness to pay and his own expected value reflects a speculative motive, the value of the right to sell the asset in the future. Moreover, uninformed or irrational collectors can interpret a surge in volume as a signal of buying interest of informed collectors and therefore decide to buy when the volume of transactions increases. Changes in volume can thus be correlated to changes in "sentiment" and predict returns (Penasse and Renneboog, 2014).

I use two proxies for volume. The most straightforward metric is the number of transactions observed each quarter. I also use Sotheby's net turnover as alternative proxy. Quarterly turnover is downloaded from CRSP and is only available from 1988Q2. This latter proxy is particularly attractive because it corresponds to value-weighted transactions, and therefore contains information about current prices. Newspapers regularly report the dollar value of items sold at auctions and it is quite likely that collectors infer current price levels from the press. I respectively denote these proxies by Volume and

Turnover. Volumes are known to be unevenly distributed within the year (as can be seen on Figure 2.2, which graphs the number of sales per months in our sample). I therefore construct Volume and Turnover series as 4-quarters sums to address seasonality, i.e. :

$$x_t = \ln(v_t + v_{t-1/4} + v_{t-1/2} + v_{t-3/4}) - \ln(v_{t-1/4} + v_{t-1/2} + v_{t-3/4} + v_{t-1})$$

where v_t denotes volume at date t. The correlation between the log-difference of the two series is 0.56.

[**Insert Figure 2.2 about here**]

To mitigate concerns that trading volume may not only capture the level of market sentiment,[9] I also use a proxy provided by the Survey of Professional Forecasters (SPF). The SPF is run by the Philadelphia Fed and provides point forecasts and expected probability distributions for inflation and output. In particular the SPF produces an "anxious index" by averaging the individual respondents' probabilities of decline in real output in the following quarter. This index has been shown to be correlated with the NBER business cycle periods of expansion and recession and is also expected to covary with macroeconomic sentiment. I therefore include this series, denoted SPF, as potential predictor.

With the exception of the SPF, all variables are returns or expressed in log-difference (see Table II). These transformations ensure that we deal with stationary data.

[**Insert Table II about here**]

IV. Do presale estimates reflect up-to-date information?

The introduction of this paper argues that price aggregation is non-trivial in a market where heterogeneous goods trade infrequently. A perfect illustration of this issue is the (relative) inaccuracy of presale estimates. Auction house experts are, by their position, in

[9]For example, if the supply for works of art is inelastic, changes in demand will affect both prices and volume (Penasse and Renneboog, 2014).

the best position to provide an accurate forecast of hammer prices and presale estimates generally provide better forecasts than econometric models based on objective hedonic characteristics. For example, in the sample where presale estimates are available, the hedonic model yields a 0.61 R-square, where a regression of hammer price on expert estimates[10] results in a 0.94 R-square. This superior performance reflects the importance of pricing factors that obviously remain unobserved to the econometrician: aesthetic quality, identity of the seller, etc.

Presale estimates may be biased, however, as experts may choose to manipulate them to influence hammer prices (Mei and Moses, 2005). Figure 2.3 charts the hedonic price index as well as the average pricing errors over the period 1995-2010. The latter are defined as

$$e_t = \frac{1}{K}\sum_{k=1}^{K} \ln P_{k,t} - \ln P^e_{k,t}$$

for each $t = 1995 \ldots 2010$, where $P^e_{k,t}$ denotes the presale estimate for item k at time t. Over the fifteen years of the sample, estimates have been on average 11% below hammer prices. This bias is highly persistent: the autocorrelation of average errors is a highly significant 0.69. This means that collectors could readily improve presale estimates by correcting for this systematic bias.

[Insert Figure 2.3 about here]

Can experts correctly aggregate market information? As can be seen on Figure 2.3, the errors present a similar trend as the price index over the period 1995-2010. Moreover, high prices seem to be associated to larger errors, suggesting that experts fail to fully anticipate price shocks.[11] To confirm this observation, I regress changes in errors on contemporaneous art returns. The regression yields an adjusted R-square of 0.25, and the slope coefficient is 0.16 with significant t-statistic of 2.35. Neither lagged changes in errors nor lagged returns seem to have predictive power of changes in errors. Experts underestimate price volatility and therefore make larger mistakes in rising than in falling markets. This finding is consistent with a behavioral explanation of short-term price

[10] The usual practice is for the auctioneer to provide a high estimate and a low estimate in an auction catalogue. I take the midpoint of the high and low estimates as forecast of the hammer price.

[11] Ashenfelter and Graddy (2011) find that unexpected price shocks are significantly correlated to the sale rate.

fluctuations, if one assumes that experts are less exuberant than collectors. The natural next step is therefore to look for the informational content of art returns.

V. The MIDAS framework

The Mixed Data Sampling (MIDAS) approach has been introduced in the econometric literature by Ghysels et al. (2004). Variables of mixed frequencies can be used in a single univariate regression model. More specifically, a MIDAS regression allows to predict a low-frequency variable with exogenous variables of higher frequency within a parsimonious and data-driven framework. Let r_t denote art returns from year $t-1$ to year t. I relate this low-frequency return to a variable $x_t^{(f)}$, which is sampled at a higher frequency than r_t. $x_t^{(f)}$ is sampled f times over the period $[t, t-1]$. In this paper f will be equal to 12 for financial variables and 4 for sentiment-related variables that are only available at quarterly frequency.

Suppose that we are interested in forecasting annual art returns using twelve lags of monthly stock returns as predictors. The forecast occurs h months before t, therefore I write monthly stock returns as $x^{(12)}_{t-(h+k-1)/f}$ where k ranges from 1 to 12. The conventional approach, in its simplest form, consists in constructing annual averages using the twelve observations of stock returns during the year: $x_{t-h/12} = (x^{(12)}_{t-h/12} + x^{(12)}_{t-(h+1)/12} + \cdots + x^{(12)}_{t-(h+11)/12})/12$ and subsequently estimate $r_t = c + \beta x_{t-h/12} + \epsilon_t$. However, a model that imposes such equal weights may suffer from inconsistent and potentially biased estimates (Andreou et al., 2010). An alternative approach is to regress r_t on each of the year's twelve monthly stock returns separately: $r_t = c + \beta_0 x^{(12)}_{t-h/12} + \beta_1 x^{(12)}_{t-(h+1)/12} + \cdots + \beta_{12} x^{(12)}_{t-(h+11)/12} + \epsilon_t$. The cost is parameter proliferation, because such a model requires to estimate 13 coefficients. MIDAS models offer a parsimonious alternative by employing a weighted time aggregation. The weights will depend of the elapsed time between sampled data and an estimated vector of hyperparameters.

Since art returns may be autocorrelated,[12] I consider a MIDAS regression augmented

[12] A number of papers have shown that past returns can help predicting future returns for collectibles (Cutler et al., 1991; Pesando, 1993; Goetzmann, 1995; Erdos and Ormos, 2010; David et al., 2013). Return autocorrelation can be explained by the "Working effect" (Goetzmann et al., 2011), the reserve-price mechanism (David et al., 2013) or collectors' under/overreaction (e.g. Erdos and Ormos 2010).

with a first-order autoregressive component:

$$r_t = c + \rho r_{t-1} + \beta \sum_{k=1}^{k^{max}} \Gamma(k,\theta) x^{(f)}_{t-(h+k-1)/f} + \epsilon_t \tag{2.4}$$

The first part of above equation is an autoregressive term that captures the information contained in the low frequency variable. The second part corresponds to the higher frequency information provided by the predictor.[13] The weights $\Gamma(k,\theta)$ control the polynomial weights that allow the frequency mixing. They are governed both by the elapsed time k and by an n-dimensional vector of hyperparameters θ. The slope coefficient β is identified via the scaling of the weights, such that they add up to one. I chose to use a two-parameter exponential Almon lag polynomial (i.e. $n = 2$):

$$\Gamma(k,\theta) = \frac{\exp(\theta_1 k + \theta_2 k^2)}{\sum_{\kappa=0}^{k^{max}} \exp(\theta_1 \kappa + \theta_2 \kappa^2)} \tag{2.5}$$

The parameters (c, ρ, β, θ) are estimated by Nonlinear Least Squares (NLS). The exponential Almon polynomial is popular in the literature because it imposes a parsimonious and reasonable restriction for which the weights are always positive. In unreported exercises, I have also experimented alternative parameterizations of the weight functions (see Ghysels et al. (2007) for a discussion of functional form choice), namely beta lag polynomial and step functions and obtained qualitatively similar results.

The parameters in Equation (2.4) depend on the forecasting horizon h, i.e. the difference between the forecast target period and the period of the last observation of the predictor. As a consequence, distinct models are estimated for different data combinations as the corresponding h varies. The case of $(h + k - 1)/f < 1$, i.e. when within year information is included in the right-hand-side, is sometimes referred as "MIDAS with leads" or "nowcasting".

Finally, the model also requires the specification of a maximum lag in the higher frequency observations. Due to the limited number of observations of art returns, I prefer to focus on relatively short-term dynamics. I use one-year information of higher frequency predictors (i.e. $k^{max} = f$), in addition to lagged art returns. For example, to

[13] With only 37 years of observations, considering several predictors simultaneously very quickly leads to parameter proliferation. I therefore estimate Equation (2.4) one parameter at a time.

predict returns 6 months ahead with monthly stock returns I regress:

$$r_t = c + \rho r_{t-1} + \beta \sum_{k=1}^{12} \Gamma(k, \theta) x^{(12)}_{t-(6+k-1)/12} + \epsilon_t$$

The above equation means that when forecasting year-end returns at the end of June, I use the last 12 observations of stock market returns (i.e. July of the previous year until June of the current year) in addition to previous year art market return.

VI. In-sample predictions

Before turning to the empirical results, it is important to consider the way in which parameter values from the estimated model (2.4) can be interpreted. The parameters θ govern the weighting applied to the predictor at each lag and cannot be given any economic interpretation. The variable of interest is therefore β and the relevant test of whether a given variable has "predictive" power is a test against the null hypothesis of $\beta = 0$. However one should be careful when interpreting such a test when the regression includes leads (i.e. $(h + k - 1)/f < 1$). When $(h + k - 1)/f \geq 1$, the horizon is at least one year so that a significant β means that the variable has predictive power. When $(h + k - 1)/f < 1$ ("nowcasting"), a significant β means that (i) a forecaster can use the variable as a proxy for unobserved within-year returns or (ii) the variable actually predicts returns, or (iii) both. Said otherwise, both correlation and predictive power can produce a significant β when the regression includes leads. For example, a within-year surge in confidence may boost art prices and volume at the same time. Both variables would induce a positive and correct forecast for year-end returns, simply because changes in volume and art returns are correlated.

Table III presents the results of the MIDAS regressions. Each column corresponds to a given regression horizon, from 3 months to one year. The first three columns therefore correspond to "nowcasting", while the fourth column considers traditional forecasting. As can be seen in Figure 2.2, a large portion of sales (40.3%) occurs in the last quarter. This implies that annual art returns should remain difficult to predict, even 3 months ahead. For each potential predictor, the table reports the value of β and of the autoregressive coefficient ρ, as well as the adjusted R^2 of the model. Heteroscedasticity and autocorrelation consistent (HAC) t-stats are reported in brackets.

[Insert Table III about here]

The first striking result is the relatively low informational content of financial variables. Albeit statistically significant, Sotheby's stock price has a modest explanatory power in the short-run. As conjectured in Section II, the explanatory power of Sotheby's stock increases with the forecast horizon.[14] The model explains 3% of art prices variance at 3-month horizon, rising to 22% on a one-year horizon. Art prices clearly react to information contained in Sotheby's price, but with a substantial lag. Surprisingly, equities have a better short-term predictive power, with a R-square of nearly 18%. The coefficient associated to lagged equity prices is only borderline significant for horizons of 3 and 9 months. Again, for a one-year forecast horizon, the statistical and economic importance of stock returns increases strongly, consistently with Goetzmann (1993) and Chanel (1995).

By contrast, the impact of trading volume, either measured by the number of sales or by Sotheby's turnover, is economically large and significant at all horizons. For horizons below six months, Volume can explain twice more art return variance than Equity. The R^2 statistic of Turnover is even higher, reflecting the price information that is not contained in Volume. As predicted by the short model of Section II volume variables progressively lose informativeness as horizon increases, while financial variables seem to gain in precision.

Finally, the survey-based "Anxious Index" has a large informational content at all horizons. Concerns of declines in real output in the following quarter are followed by negative returns the same year. SPF has higher predictive power than financial variables for short-term horizons. As stressed earlier, a plausible explanation of the joint role of volume and SPF is the role of market sentiment. For example, Renneboog and Spaenjers (2013) find that lagged equity returns lose statistical significance when controlling for contemporaneous market sentiment and conclude that time-varying optimism about art investment impacts art pricing.

[14] I only discuss R-squareds, because the right-hand scale variables are weighted averages of monthly or quarterly variables so that the value of β is difficult to interpret.

VII. Out-of-sample predictions

In this section, I evaluate the out-of-sample performance of MIDAS regressions with respect to two benchmark investors. The first consists of an uninformed investor who is skeptical about art returns predictability. Such an investor would predict returns to be equal to their historical average, i.e. use Equation (2.2). The second benchmark investor is an informed agent who observes the interim index described by Equation (2.3).

I compare the above forecasts from Equations (2.2) and (2.3) with the predictions from the MIDAS regressions. I consider the same forecast horizons h as in the previous section. I first estimate the MIDAS regression model (V) on an observation period consisting of the sample from 1974 to 1998 (1989 to 2002 for Sotheby's stock price and Turnover, which are only available from 1988). Then for each year after 1998 (2002), I re-estimate the model in real time using an expanding window up to 2009. The out-of-sample forecasts write as:

$$\hat{r}_{t|t-h/f,x} = \hat{c} + \hat{\rho} r_{t-1} + \hat{\beta} \sum_{k=1}^{k^{max}} \Gamma(k, \hat{\theta}) x^{(f)}_{t-(h+k-1)/f} \tag{2.6}$$

where the parameters $(\hat{c}, \hat{\rho}, \hat{\beta}, \hat{\theta})$ are estimated dynamically at time $t - h/f$. Since the forecasts obtained by the tracking index are expected to be imprecise, one would like to see if the latter could be improved using MIDAS forecasts. Model averaging techniques often increase the precision of the forecasts relative to those of individual models. Therefore, in addition to these out-of-sample forecasts obtained from single predictor regressions, I construct forecast combinations from the single-variable models:

$$\hat{r}_{t|t-h/f,\text{avg}} = \sum_{m=1}^{M} w_{m,t-h/m} \hat{r}_{t|t-h/m,x=m},$$

where $\hat{r}_{t|t-h/m,x=m}$ is the prediction associated to model m and M is the number of models considered. I use a uniform average, i.e. $w_{m,t-h/m} = 1/M$ and a recursive weighed-scheme:

$$w_{m,t-1} = \frac{\exp\left(-\frac{1}{2}\text{BIC}_{m,t-1}\right)}{\sum_{l=1}^{M} \exp\left(-\frac{1}{2}\text{BIC}_{l,t-1}\right)}$$

where $\text{BIC}_{m,t-1}$ is the Bayesian Information Criterion of model m observed at time $t-1$. I consider forecast combinations including and excluding the interim index.

In order to assess forecast accuracy, I compute the ratios of Root Mean Squared Forecasting Errors (RMSFEs) of the model-based forecasts and of constant annual returns forecasts (2.2):

$$r^{(h)} = \frac{\text{RMSFE}^{(h)}_{\text{model}}}{\text{RMSFE}_{\text{constant returns}}}.$$

Therefore a ratio $r^{(h)} < 1$ implies that MIDAS forecasts are able to beat the historical average.

Table IV shows the forecast ratios of these models. The most striking result is the superior performance of the tracking index, which ratio of RMSFE is always below one for all horizons below one year. It is exactly equal to one for a one-year horizon when the two forecasts (2.2) and (2.3) are actually identical. The ratio is by far lower than alternative models, including forecast combinations. The performances of the latter are at best equivalent to the index performances, when a weight of approximately one is applied to the tracking index forecasts.

[**Insert Table IV about here**]

While Volume is able to do better than the unconditional mean up to six months, return forecastability does not extend to Turnover. Forecast combinations of MIDAS models, reported on the last two lines, all fail to beat the historical average.

The inability of financial variables to forecast returns in real time is illustrated on Figure 2.4. Figure 2.4 plots annual updates of the price index, together with forecasts obtained from the tracking index and from the MIDAS model using either Equity or Volume as predictor. The tracking index does a good job, in particular in anticipating large price moves. Such moves, e.g. the 2004 surge in price and the 2008 drop, are already priced in from the first quarter. Equity forecasts, by contrast, are much more hazardous. In particular, the model wrongly predicts a large drop following the collapse of the internet bubble and fails to forecast the 2006-2008 boom. Volume forecasts seem to track art prices quite well. They are much less volatile than Equity forecasts and, as a result, the model never makes large mistakes. Moreover, the model correctly predicts large price changes, in contrast to Equity forecasts.

[**Insert Figure 2.4 about here**]

As discussed in the previous section, for horizons less than one year, one cannot distinguish between predictability and correlation. The last column of Table IV, which presents the forecast for a one-year horizon, thus highlights the absence of out-of-sample predictability. This lack of predictability may be caused by the short size of our sample, but it is robust to alternative initial observation periods.[15] Such unpredictability is consistent with similar evidence on the stock market (see, e.g. Welch and Goyal 2008), but is novel to the art market.

VIII. Conclusion

In this article, I assess the ability of macro-financial and proxies for market sentiment in forecasting year-end art prices. I use the MIDAS regression framework, which allows to explain a low-frequency variable by exogenous variables of higher frequency within a parsimonious framework. The objective is twofold. From a theoretical point of view, this paper measures how information from higher frequency markets diffuses into art prices. From a practical point of view, it assesses the capacity of these variables to improve out-of-sample forecasts of art prices.

It takes about six months for art prices to incorporate information contained in the price of Sotheby's stock. While conventional wisdom views the art market as increasingly linked to financial markets, the short-term financial content of art returns is relatively poor, but increases with the investment horizon. By contrast, for horizons below six months, variables related to market sentiment have a much larger explanatory power. Out-of-sample, art returns are largely unpredictable. Only forecasts exploiting volume information can improve predictions of investors, provided they do not have access to interim prices.

These findings indicate that the art market incorporates information about macroeconomic fundamentals with a substantial lag, because information diffuses slowly across markets. A plausible explanation is market segmentation. The proxies for market sentiment are likely to better reflect art market participants' expectations and hence have better forecasting power. Slow information diffusion can also arise because art market

[15]The unreported results are available upon request.

participants never observe aggregate art prices in real time, which further complicates information aggregation.

Bibliography

Andreou, Elena, Eric Ghysels, and Andros Kourtellos, 2010, Regression models with mixed sampling frequencies, *Journal of Econometrics* 158, 246–261.

Ashenfelter, Orley, and Kathryn Graddy, 2003, Auctions and the Price of Art, *Journal of Economic Literature* 41, 763–786.

Ashenfelter, Orley, and Kathryn Graddy, 2011, Sale rates and price movements in art auctions, *American Economic Review* 101, 212–216.

Baker, Malcolm, and Jeremy C. Stein, 2004, Market liquidity as a sentiment indicator, *Journal of Financial Markets* 7, 271–299.

Brennan, Michael J., Narasimhan Jegadeesh, and Bhaskaran Swaminathan, 1993, Investment analysis and the adjustment of stock prices to common information, *Review of Financial Studies* 6, 799–824.

Burton, Benjamin J., and Joyce P. Jacobsen, 1999, Measuring returns on investments in collectibles, *Journal of Economic Perspectives* 13, 193–212.

Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay, 1997, *The econometrics of financial markets* (Princeton University Press).

Chanel, Olivier, 1995, Is art market behaviour predictable?, *European Economic Review* 39, 519–527.

Cohen, Lauren, and Andrea Frazzini, 2008, Economic Links and Predictable Returns, *Journal of Finance* 63, 1977–2011.

David, Géraldine, Kim Oosterlinck, and Ariane Szafarz, 2013, Art market inefficiency, *Economics Letters* 121, 23–25.

Erdos, Péter, and Mihály Ormos, 2010, Random walk theory and the weak-form efficiency of the US art auction prices, *Journal of Banking & Finance* 34, 1062–1076.

Frank, Robert, Art is a bubble: Here's how to short it, 2014. CNBC.

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

Ghysels, Eric, Pedro Santa-Clara, and Rossen Valkanov, 2004, The MIDAS touch: Mixed data sampling regression models, *CIRANO Working Papers* 2004s-20.

Ghysels, Eric, Arthur Sinko, and Rossen Valkanov, 2007, MIDAS Regressions: Further Results and New Directions, *Econometric Reviews* 26, 53–90.

Ginsburgh, Victor, J. Mei, and M. Moses, 2006, The Computation of Prices Indices, *Handbook on the Economics of Art and Culture* 1, 947–979.

Goetzmann, William N., 1993, Accounting for taste: Art and the financial markets over three centuries, *American Economic Review* 83, 1370–1376.

Goetzmann, William N., Elena Mamonova, and Christophe Spaenjers, 2014, The Economics of Aesthetics and Three Centuries of Art Price Records, *NBER Working Paper* 20440.

Goetzmann, William N., Luc Renneboog, and Christophe Spaenjers, 2011, Art and money, *American Economic Review* 101, 222–226.

Hiraki, Takato, Akitoshi Ito, Darius A. Spieth, and Naoya Takezawa, 2009, How did Japanese investments influence international art prices?, *Journal of Financial and Quantitative Analysis* 44, 1489.

Hong, Harrison, Terence Lim, and Jeremy C. Stein, 2000, Bad News Travels Slowly: Size, Analyst Coverage, and the Profitability of Momentum Strategies, *Journal of Finance* 55, 265–295.

Hong, Harrison, and Jeremy C. Stein, 1999, A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets, *Journal of Finance* 54, 2143–2184.

Hong, Harrison, Walter Torous, and Rossen Valkanov, 2007, Do industries lead stock markets?, *Journal of Financial Economics* 83, 367–396.

Hou, Kewei, 2007, Industry information diffusion and the lead-lag effect in stock returns, *Review of Financial Studies* 20, 1113–1138.

Huberman, Gur, and Tomer Regev, 2001, Contagious Speculation and a Cure for Cancer: A Nonevent that Made Stock Prices Soar, *Journal of Finance* 56, 387–396.

Korteweg, Arthur, Roman Kräussl, and Patrick Verwijmeren, 2013, Does it pay to invest in art? A selection-corrected returns perspective, *LSF Research Working Paper Series* 13–07.

Lo, Andrew W., and Archie Craig MacKinlay, 1990, When are contrarian profits due to stock market overreaction?, *Review of Financial Studies* 3, 175–205.

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

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

Menzly, Lior, and Oguzhan Ozbas, 2010, Market Segmentation and Crosspredictability of Returns, *Journal of Finance* 65, 1555–1580.

Miller, Edward M., 1977, Risk, uncertainty, and divergence of opinion, *Journal of Finance* 50, 1013.

Penasse, Julien, and Luc Renneboog, 2014, *Bubbles and Trading Frenzies: Evidence from the Art Market*, Working paper, Tilburg University.

Penasse, Julien, Luc Renneboog, and Christophe Spaenjers, 2014, Sentiment and art prices, *Economics Letters* 122, 432–434.

Pollet, Joshua A., and Stefano DellaVigna, 2007, Demographics and industry returns, *American Economic Review* 97.

Pownall, Rachel A.J., Stephen Satchell, and Nandini Srivastava, 2013, *A random walk through Mayfair : dynamic models of UK art market prices and their dependence on UK equity prices*, Working paper, Maastricht University.

Rapach, David E., Jack K. Strauss, and Guofu Zhou, 2013, International Stock Return Predictability: What Is the Role of the United States?, *Journal of Finance* 68, 1633–1662.

Renneboog, Luc, and Christophe Spaenjers, 2013, Buying beauty: on prices and returns in the art market, *Management Science* 59, 36–53.

Scheinkman, Jose A., and Wei Xiong, 2003, Overconfidence and speculative bubbles, *Journal of Political Economy* 111, 1183–1220.

Welch, Ivo, and Amit Goyal, 2008, A Comprehensive Look at The Empirical Performance of Equity Premium Prediction, *Review of Financial Studies* 21, 1455–1508.

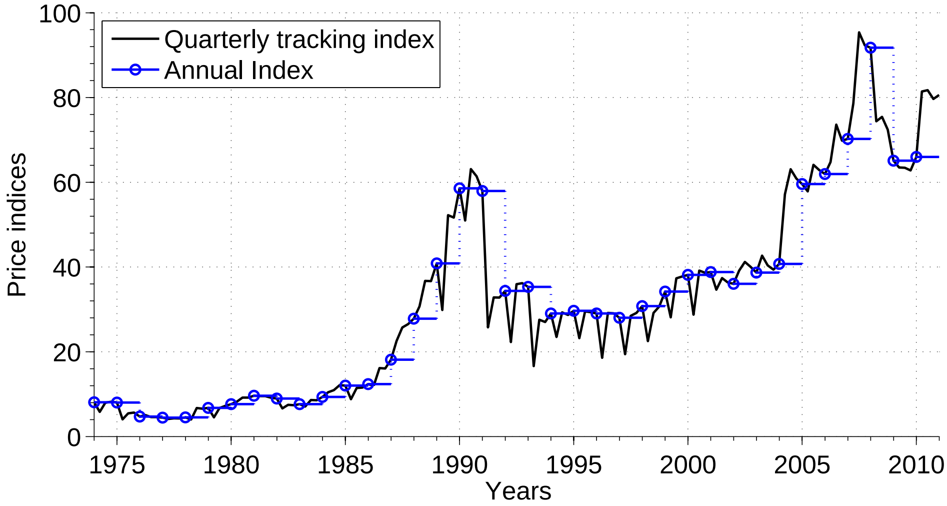

Figure 2.1: Art Indices

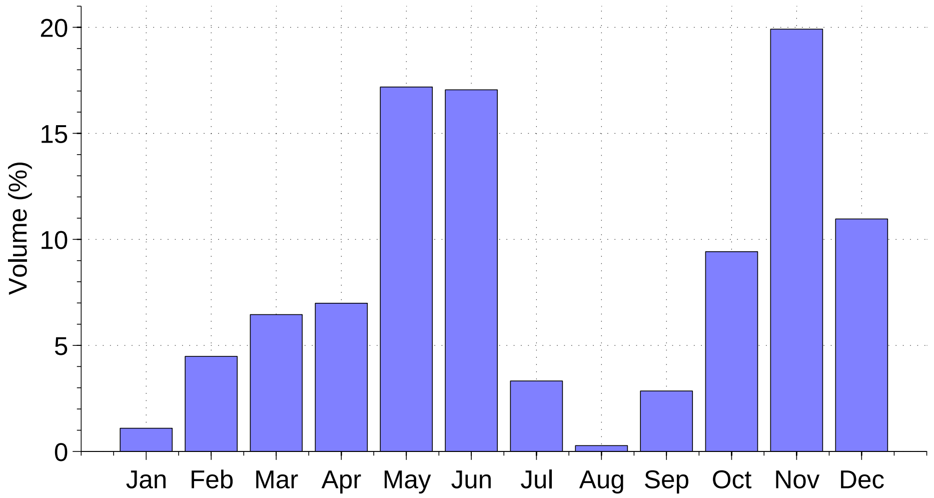

Figure 2.2: Within-year distribution of sales

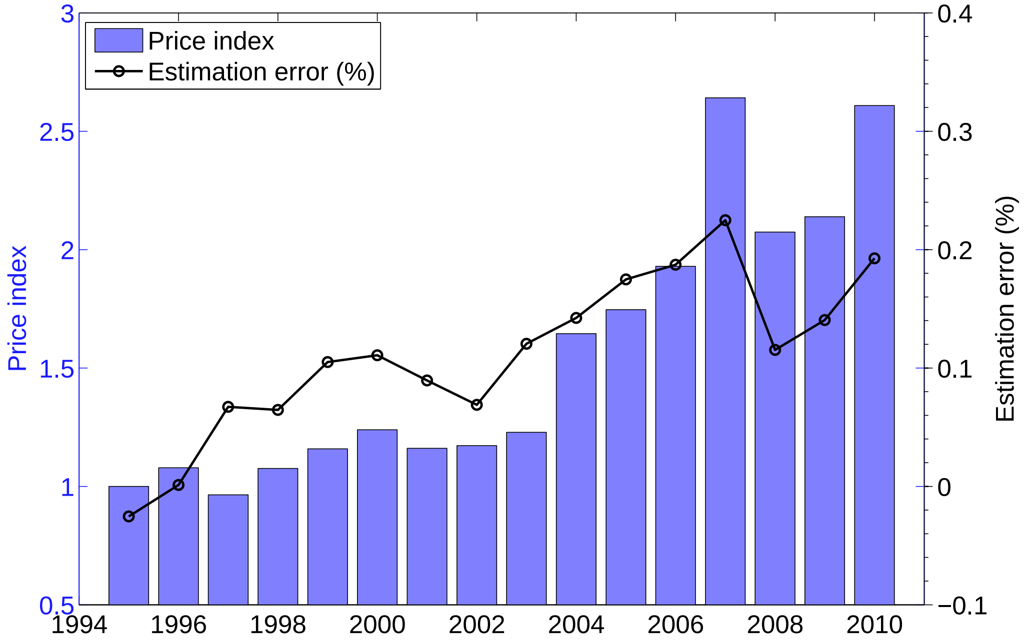

Figure 2.3: Price index and average estimation error
Notes: this figure charts the hedonic price index and average estimation error over the period 1995-2010. Estimation errors are computed as the log difference between the outcome of each sale and its presale estimate.

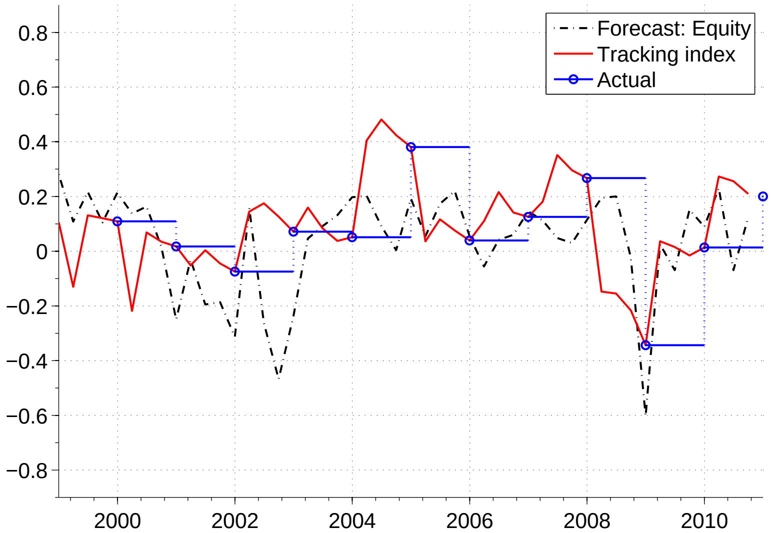

(a) Equity

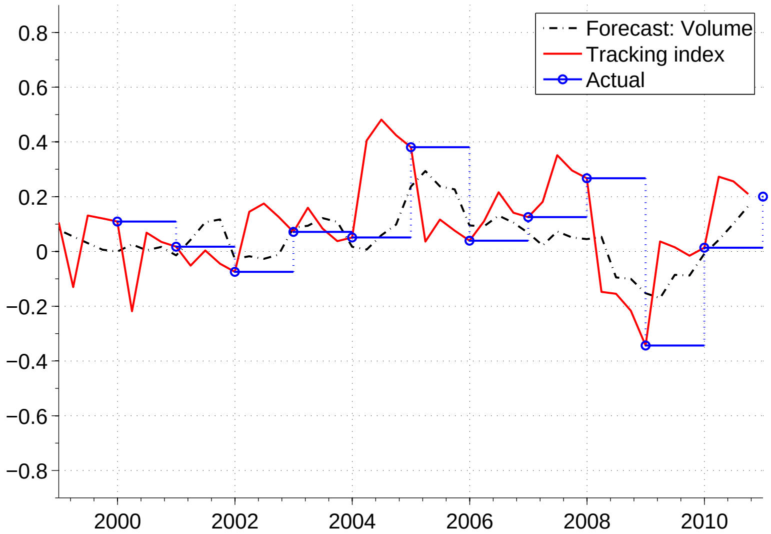

(b) Volume

Figure 2.4: Actual year-end returns and real-time forecasts: illustration for Equity and Volume

This figure plots actual realizations of the price index and forecasts obtained from the tracking index (using Equation (2.3)) and forecast from MIDAS regressions using either Equity or Volume as proxy for art returns. Each dot corresponds to a new realization of the price index, which is updated annually. The true realizations of the price index are compared to forecasts performed from one year to three months ahead, and appear on the left of each dot.

Table I: Hedonic regression results

Variable	Coefficient	t-stat	N
Time		[Included]	
Artists		[Included]	
Technique			
Oil	1.88	75.88	36,073
Other painting	0.72	22.36	3,018
Ohotograph	−0.55	−7.61	348
Reproductible	−1.19	−48.8	36,826
Sculpture	−0.22	−6.87	3,369
Works on paper	0.1	4.43	59,208
Size			
Height	0.006	35.41	
Width	0.004	26.37	
Height^2	−0.000	−38.3	
Width^2	−0.000	−22.69	
Other Characteristics			
Signed	0.72	69.43	74575
Series	0.81	79.58	8431
Dated	−0.1	−3.77	74280
Sale			
Christie's, London	0.36	45.89	16,845
Christie's, New York	−1.11	−62.68	19,440
Christie's, other	0.33	38.9	3,905
Sotheby's, London	0.68	61.14	23,351
Sotheby's, New York	0.81	76.91	21,661
Sotheby's, other	−0.18	−8.78	2,288
c	4.79	38.3	141,638
Adj R^2	0.60		

This table presents the hedonic regression results based on the full sample (1954-2010). The model, presented in Equation (2.1), is estimated using OLS. The Time and Artist dummies identify the year of the sale and the name of the artist. The dummies Oil, Other painting, Photograph, Reproducible, Sculpture and Works on paper indicate the technique. The Other Characteristics dummies Signed, Series and Dated take the value one if a work is signed, part of a series, or is dated, respectively.

Table II: Description of variables

	Description	Frequency	Transformation
Art index	Annual art index updated in real time, 1973-2010	Annual	log-return
Tracking index	Annual index tracking updates during the year, updated in real time on a quarterly basis, 1973Q4-2010Q4	Quarterly	log-return
Sotheby's	Sotheby's stock (Bloomberg)	Monthly	log-return
Equity	MSCI World	Monthly	log-return
Bond	10-year US government benchmark (Datastream)	Monthly	log-return (constant maturity)
Volume	1-year rolling number of transactions	Quarterly	log-difference
Turnover	1-year rolling Sotheby's turnover (CRSP)	Quarterly	log-difference
SPF	Survey of Professional Forecasters: probability of decline in real GDP in the following quarter (Philadelphia Fed) Quarterly	Quarterly	-

Table III: MIDAS predictive regressions

		Art returns, m months ahead			
		$m = 3$	$m = 6$	$m = 9$	$m = 12$
Sotheby's	β	0.4098	0.7620	2.3814	2.1995
		(1.7085*)	(1.8329*)	(2.4742**)	(2.6536**)
	ρ	−0.0159	−0.0210	−0.2364	−0.2501
		(−0.0802)	(−0.1072)	(−1.1150)	(−1.2372)
	R^2	0.0321	0.0504	0.1839	0.2213
Equity	β	3.4720	2.8693	2.0687	7.3428
		(1.8387*)	(2.6823**)	(1.9651*)	(3.4501***)
	ρ	0.2981	0.2342	0.2522	0.1068
		(1.9722*)	(1.5624)	(1.6152)	(0.7167)
	R^2	0.1790	0.1734	0.1025	0.2596
Bond	β	−0.6636	1.0234	0.1656	−0.6388
		(−0.5721)	(0.7615)	(0.1820)	(−0.7950)
	ρ	0.2259	0.2077	0.2461	0.2265
		(1.3534)	(1.2227)	(1.4916)	(1.3212)
	R^2	0.0171	0.0344	0.0045	0.0259
Volume	β	3.5368	3.3821	1.5959	1.3360
		(4.0982***)	(4.0343***)	(2.3875**)	(2.4616**)
	ρ	0.2586	0.2067	0.2146	0.1848
		(1.9123*)	(1.4651)	(1.3258)	(1.1739)
	R^2	0.3361	0.3319	0.1491	0.0851
Turnover	β	1.6636	1.6439	2.2070	1.2806
		(3.1179***)	(3.2566***)	(4.4369***)	(2.7000**)
	ρ	−0.2076	−0.2042	−0.2768	−0.4346
		(−1.3472)	(−1.3243)	(−1.5374)	(−1.7462*)
	R^2	0.4810	0.4819	0.4148	0.1931
SPF	β	−0.0060	−0.0062	−0.0062	−0.0045
		(−3.0934***)	(−3.3209***)	(−3.1391***)	(−2.2327**)
	ρ	0.1973	0.1988	0.1947	0.1806
		(1.3502)	(1.3767)	(1.3587)	(1.1690)
	R^2	0.2358	0.2482	0.2603	0.1756

This table presents the estimation results of the autoregressive MIDAS model in Equation (2.4). The dependent variable is annual art market returns and the regression is performed using monthly or quarterly variables with horizon $h = m/12$ (see Table II for a description of the variables). The specification involves an AR(1) autoregressive term and an exponential Almon lag MIDAS polynomial. I use one year of observations, i.e. 4 lags of quarterly variables and 12 lags of monthly variables. The model is estimated by NLLS. For the sake of brevity, I only present the estimates of the (first-order) autoregressive parameter ρ, and the slope parameter β. The terms in brackets are the associated (HAC corrected) t-stats. The R^2 statistics are adjusted for the number of explanatory terms.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table IV: Out-of-sample forecast performances

	Art returns, m months ahead			
	$m = 3$	$m = 6$	$m = 9$	$m = 12$
Tracking index	0.27	0.5	0.71	1.00
MIDAS				
Sotheby's stock*	1.29	0.99	1.83	1.49
Equity	1.36	1.38	1.06	1.67
Bond	1.19	1.32	1.14	1.18
Volume	0.90	0.93	1.18	1.11
Turnover*	1.08	1.12	1.02	1.33
SPF	1.25	1.25	1.24	1.25
Forecast combinations : MIDAS + Tracking index				
Uniform	0.87	0.92	0.93	1.13
BIC	0.27	0.5	1.24	1.64
Forecast combinations : MIDAS only				
Uniform	1.03	1.08	1.07	1.11
BIC	1.05	1.2	1.24	1.65

This table reports the ratios of RMSFEs obtained when forecasting art returns using the tracking index or MIDAS regressions and of RMSFEs obtained from predictions based on the constant return benchmark. A ratio below one indicates predictive ability with respect to the historical average. The forecasts based on the tracking index and on MIDAS regressions are also combined through uniform and BIC-weighted averages.
* The forecasts are performed on the out-of-sample period (1998-2010) after an observation period from 1974 to 1998 (with the exception of Sotheby's stock and Turnover, which observation period is 1988 to 2002 due to data limitation).

Chapter 3

Bubbles and Trading Frenzies: Evidence from the Art Market[1]

Chapter Abstract

The art market is subject to frequent booms and busts in both prices and volume, which are difficult to reconcile with models where agents are rational and hold homogenous beliefs. This paper shows that (i) volume is mainly driven by speculative transactions; (ii) positive price-volume correlation is pervasive across art movements, and is larger for the most volatile segments of the art market; (iii) volume predicts negative long-term returns, a relation that is statistically and economically large. Overall, our evidence supports the bubble model of Scheinkman and Xiong (2003), which predicts that speculative trading can generate significant price bubbles, even if trading costs are huge and leverage is impossible.

[1]We are grateful to Fabio Braggion, Edouard Challe, Guillaume Chevillon, Pierre Collin-Dufresne, Gabriel Desgranges, Joost Driessen, Andreas Heinen, Frank de Jong, Hao Liang, Benjamin Mandel, Kim Oosterlinck, Christophe Pérignon, Olivier Scaillet, Christophe Spaenjers and participants of the French Finance Association Conference (Strasbourg), the History of Economics as Culture 4th Annual Workshop (Cergy) and seminars at THEMA, ESSEC and Tilburg University for helpful comments and suggestions. A first version of this paper was entitled *A Beauty Contest? Disagreement in the Art Market.*

I. Introduction

As a collector, I trade all the time, it's the capitalist in me.[2]

In January 1989, Financial Times journalist Robin Duthy (1989) wrote: "The art market today is in a sound state, but the danger is that the long run of sparkling results for paintings by Monet, Van Gogh, and other household names will create an illusion that all art is safely strapped in on some kind of magic escalator." In the five-year period from 1985 to 1989, art prices had grown by 164%, in a context of record sales and apparent overoptimism. In the subsequent two years, real prices bounced back to their 1986 levels. A similar run up in prices, and subsequent collapse, was to a lesser extent reproduced in 2002-2008, and many are those who see another "bubble" in current prices.

Speaking of "bubbles" requires a definition of fundamental value, which is challenging when applied to art. As art prices soar, art dealers, auctioneers, and art gallery sales people often emphasize the resell value, while after a bust, they tend to comfort collectors by stating that pleasure is the best dividend when investing in art. For economists, works of art differ from traditional assets or durable goods in that they yield a non-pecuniary aesthetic or utility dividend. This utility dividend can be seen as the rent one would be willing to pay to own this work of art over a given time frame. It can reflect aesthetic pleasure but also has the ability to signal its owner's wealth. The value of this dividend is of course unobservable and is likely to vary tremendously across art collectors. As pointed by Lovo and Spaenjers (2014), however, the auction market introduces a common-value element into prices. The price of a work of art should therefore equal the present value of future (private) utility dividends over one's expected holding period, plus the expected (market) resell value, i.e. the discount rate model could be use to price art. We can readily propose a definition of fundamental value. Assuming the last bidder on a work of art is the one who values it the most, the fundamental value of an artwork is his own private value. Hence, a "bubble" corresponds to a market where agents are willing to pay more than their private value, because they expect to resell later at a larger price.[3]

[2]Quoted in the *Wall Street Journal* (Peers, 2008).

[3]Two years before the 1990 bust, the following quote appeared in the *New York Times*: "Paintings that used to sell for $400,000 are now going for $4 million to $5 million. [...] And when you pay those

Viewing art as an asset helps understand the possible sources of art price fluctuations. The most straightforward explanation is that the utility dividends fluctuate over time as they depend on buyers' willingness to pay for art, which in turn depends on their preferences and wealth[4] (Mandel, 2009). In order to explain art prices volatility, preferences regarding art and culture as a whole would have to fluctuate dramatically. Even if fads can temporarily emerge for some specific artists or school of art (Penasse et al., 2014), the previous literature has shown that tastes tend to be very stable, even in the long run (Ginsburgh and Weyers, 2008; Graddy, 2014).[5] Alternatively, the utility dividend can oscillate because people's wealth fluctuates over time. The literature has provided evidence supporting this idea, which we denote as the *luxury consumption* hypothesis. For example Goetzmann et al. (2011) find cointegrating relationships between top incomes and art prices. Finally, art's fundamental value can fluctuate because the discount rate, i.e. the risk premium associated to holding works of art changes over time.[6]

A salient feature of art price booms is that they are accompanied, and sometimes preceded, by large volumes. Figure 3.1 shows that the total number of transactions rose 45% from 1985 until its peak in 1989, and many segments of the market reached much higher levels. For example, during the same period, the prices and volume of Pop artists respectively rose 354% and 167%; more works by Andy Warhol where sold in 1989 than in the four previous years combined. The positive correlation between the art prices and volume corroborates a similar observation about many historical bubbles, such as the South Sea Bubble, or more recently the Internet bubble in the late 1990s.[7] Moreover, the share of short-term transactions, identified as purchases that were resold within the next year, rose from 10% to almost 20%. This is a remarkable increase, knowing that the

prices, you're an investor. You'll see the paintings bought at these sales come up for auction again in several years" (Glueck, 1988).

[4]Assuming that art consumption increases with wealth, i.e. that works of art are luxury goods.

[5]Citing Stigler and Becker (1977), LeRoy (2004) further argues "against relying on assumed preference shifts to explain price fluctuations, especially when there exist alternative explanations that do not appeal to preference shifts."

[6]Fluctuations in the discount rate is the most straightforward explanation of asset prices fluctuation within efficient markets (Cochrane, 2011).

[7]See e.g. Cochrane (2003); Ofek and Richardson (2003). Xiong (2013) notes that classical economists such as Adam Smith, John Stuart Mill, Knut Wicksell, and Irving Fischer readily proposed the concept of "overtrading," the process whereby euphoric investors buy assets solely in anticipation of future capital gains (Kindleberger, 1978). The first historical bubbles where readily characterized by trading frenzies. For example, Carlos et al. (2006) show that turnover in the shares of the Bank of England, the East India Company, and the Royal African Company increased dramatically during the South Sea Bubble of 1720.

transaction costs in auction markets are minimally 25% of the hammer price. Interestingly, the relation between prices and volume is not confined to a few episodes or markets. Price increases generally coincide with rises in volume: between 1976 and 2006, the correlation between changes in art prices and changes in art volume was as high as 54%.

[**Insert Figure 3.1 about here**]

What can explain such trading frenzies? In traditional asset pricing models agents cannot trade when they share identical prior beliefs (Tirole, 1982; Milgrom and Stokey, 1982), they rather trade to consume, i.e. because their preferences, or wealth, differ. The rationality assumption rules out any form of speculative trading, because in such models agents agree on economic fundamentals and therefore cannot expect to make a profit by reselling later. Several seminal papers, in contrast, emphasize the role of speculation on price formation when agents hold heterogeneous beliefs or priors (Miller, 1977; Harris and Raviv, 1993). For example, agents can trade because they are overconfident about their own trading abilities (Scheinkman and Xiong, 2003), because they suffer from confirmatory bias (Pouget et al., 2014), or because they are trying to infer what others are thinking (Biais and Bossaerts, 1998). These models suggest that "market sentiment" or differences of opinion can push prices above fundamentals. Their arguments hinge on the assumption that short-sale constraints prevent arbitrageurs from pulling back prices to fundamentals (Miller, 1977; Baker and Stein, 2004). When prices are high, pessimists would like to short sell, but instead simply stay out of the market or sell to optimists at inflated prices. Moreover, optimists may be willing to pay higher prices than their own valuations, because they expect to resell to even more optimistic investors in the future (Harrison and Kreps, 1978; Scheinkman and Xiong, 2003). The difference between their willingness to pay and their own optimistic valuation is the price of the option to resell the asset in the future. The price of the resale option imparts a stationary bubble component in asset prices, and can explain price fluctuations unrelated to macroeconomic fundamentals. This mechanism is particularly appealing in explaining art price fluctuations, because in art markets short selling is not possible and, absent a rental market, the only possibility to make a profit is by reselling at a higher price.

It is important to stress that speculative trading is not the only cause of price-volume correlation.[8] Price-volume correlation can also arise when the market is subject to supply or demand shocks, and if supply is less elastic than demand (or reciprocally). A prominent example of demand shock is the Japanese stock and real estate "bubble" of the late 1980s. Hiraki et al. (2009) document how Japanese collectors entered the market chasing French Impressionist art, pushing both prices and volumes up. They argue that their findings reflect luxury consumption, consistent with predictions of consumption capital asset pricing models (Aït-Sahalia et al., 2004). Lovo and Spaenjers (2014) present a dynamic auction model where transaction prices and voluntary trading volume increase when the economy enters an expansion, and decrease when a recession commences. A crucial difference between the two former explanations of trading volume is that the latter is silent about who trades and about the kind of goods that are traded. Further, resale option theory states that a high volume should be associated to prices above fundamental value, so that a high trading volume should predict negative returns.

In this paper, we use a comprehensive data set of nearly 1.1 million auction sales to study asset price fluctuations and trading in the fine art auction market. First, we intend to analyze what drives trading volume. The previous literature, which has extensively studied the demand for works of art, has remained silent on the informational content of trading volume. We have presented two competing theories — the speculative trading theory and the luxury consumption theory — that can explain a correlation between prices and volume. Consonant with the *speculative trading* hypothesis, we find that the share of very short-term transactions, the sales rate, and the share of the riskier art movements increase when volume increases. While prices increase with top incomes, as predicted by the *luxury consumption* hypothesis, the contemporaneous relation of top income augmentation with volume is positive but insignificant.

Second, we find that the positive contemporaneous price-volume correlation is robust and pervasive across art movements and that the riskier (respectively, safer) artists tend to exhibit higher (lower) price-volume correlation. Perhaps unsurprisingly, riskier art (Pop, Abstract Expressionism, Minimalism and Contemporary art) belongs to the second half of twentieth century, while safer art (Romanticism, Baroque, Rococo) ended no later

[8] To avoid making this paper too cumbersome, we will liberally speak of price-volume correlation. What we mean is correlation between log-differences.

than the first half of the nineteenth century.

While supportive of the *speculative trading* hypothesis, these findings are not yet conclusive. Our third objective is thus to study whether a high volume coincides with overpricing, which is the most important prediction of the resale option theory (Hong and Stein, 2007). Although the fundamental value of art is unobservable, a clear test of overpricing is that volume negatively predicts returns, while controlling for potential changes in fundamental value. Crucially, our dataset contains more than 20,000 pairs of transactions where identical items have been identified at the time of purchase and subsequent resale. This enables us to test directly the overpricing prediction of the resale option theory. In order to control for changes in fundamental value, we then turn to the classic capital asset pricing model (CAPM) to express art excess returns in terms of systematic risk, and additionally control for changes in artist fame and changes in volume. A one standard deviation increase in volume will on average increase future excess returns by 15.1% over the holding period, or 2.6% per year. This long-term effect of volume on art returns is much larger than the effect of stock returns (5.5% on average) and the effect of taste (6% on average). Importantly, this relation is robust across time, which means our results are not driven by the 1990s boom.

Our paper contributes to a number of strands of the literature.

First, we provide evidence supporting a resale option theory and, in particular, the bubble model of Scheinkman and Xiong (2003). The previous literature has provided empirical evidence related to events limited in time, such as the Chinese warrant bubble (Xiong and Yu, 2011) and the Chinese A-B share premia (Mei et al., 2009). Palfrey and Wang (2012) also find evidence of speculative overpricing in laboratory-controlled asset markets. Our finding is novel, in that overpricing is not driven by a single event. Speculation occurs in spite of huge transaction costs, as conjectured by Scheinkman and Xiong (2003). We directly relate the large fluctuations in art prices to a stationary overpricing component. We argue that this overpricing component, which is proxied by trading volume, induces predictability in art returns.[9] Bubbles in the art market are

[9]This stands in marked contrast to the purely rational view that argues that changes in the discount factor induces return predictability (see Cochrane (2011) for an overview of this literature). A recent paper by Cujean and Hasler (2014) also argues that disagreement generates predictability over the business cycle.

unique in that they can start in the absence of large uncertainty or innovation[10] and are not driven by excess credit or leverage (Stein, 1995; Geanakoplos, 2010).

Second, our data enables us to examine the empirical relation between prices and volume. While this relation has been extensively examined in other asset markets, including the housing market,[11] virtually no research exists on markets of collectibles. Ashenfelter and Graddy (2011) study sales rates at art auctions, but not volume per se. Bai et al. (2013) examine volume through the lens of international trade. Our set-up is complementary to that of Korteweg et al. (2013) and Lovo and Spaenjers (2014), who examine how changes in market values correlate with the likelihood of trading for individual artworks.

Third, our paper extends the understanding of the drivers of "emotional asset" prices. Previous research has insisted on the role of wealth, proxied by equity returns and changes in the income distribution (Goetzmann et al., 2011; Hiraki et al., 2009). Interestingly, Hiraki et al. (2009) explains the 1990 art price bubble in fundamental terms: luxury consumption by Japanese art collectors pushed international art prices up until the art bubble burst as a direct consequence of the collapse of the Japanese real estate market. We provide an alternative interpretation emphasizing speculative dynamics, which has been largely overlooked by the literature. Penasse et al. (2014) use survey data to show that optimism about individual contemporary artists has predictive power of short-run art returns, in line with the idea that fads affect the prices of individual artists. In a similar vein, Renneboog and Spaenjers (2013) construct an aggregate art market sentiment indicator, based on sales volume, buy-in rates, and the tone of press reports on the art market, which covaries with art prices. Still, that study does not address a price-volume correlation nor provides evidence of return predictability.[12]

Fourth, we shed more light on the behavioral anomalies that characterize the auction market. Mei and Moses (2005) show that high estimates at the time of purchase are

[10]Innovation and uncertainty is an inherent element of asset prices bubbles (see e.g. Xiong (2013)), which complicates the identification of bubbles even ex post. See, e.g. Pastor and Veronesi (2006) for a rational explanation of the 1990 internet bubble.

[11]See Genesove and Mayer (2001), Clayton et al. (2008); see also Piazzesi and Schneider (2009) and Favara and Song (2013) for arguments in terms of overpricing.

[12]Two recent papers also investigate the short-term dynamics of art prices. Pownall et al. (2013) employ a regime switching model to describe the dynamics of art prices using a threshold variable that drives prices into possibly locally explosive regimes. Kräussl et al. (2014) use a right-tailed unit root test with forward recursive regressions to detect explosive behaviors in the prices of four different art market segments.

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. De Silva et al. (2012) show that investors' emotional state (in their paper influenced by the weather at the time of the auction) can affect price formation of paintings with a relative high private value. Pesando and Shum (2007) present anecdotal evidence of "irrational exuberance" in the prices realized at the 1997 sale of the collection of Victor and Sally Ganz at Christie's in New York. They argue that the buyers of five Picasso prints probably overpaid, as evidenced by the dramatically lower prices realized by these prints in their subsequent appearances at auction. Our results suggest that this tendency of collectors to overpay is significant and pervasive.

The remainder of this paper is structured as follows. We present our dataset in Section II. Section III provides evidence on what drives volume. Our core results are presented in Sections IV and V, where we study price-volume correlation across art movements and show that volume has long-term predictive power. Section VI discusses the interpretation of our empirical findings and Section VII concludes.

II. Data

This paper uses the historical data set constructed by Renneboog and Spaenjers (2013), which comprises information on more than one million transactions of art at auction over the period from 1957 until 2007. The dataset initially overweighs the London sales, but as of the middle of the 1970s, the coverage consists of all major auction houses around the world. We thus concentrate our analysis of price and volume on the 1976-2006 period.[13] The sales concern oil/acrylic paintings and works on paper (water colors, gouaches, etchings, prints) by more than 10,000 artists. Almost half of the artists are classified into one or more of the following movements: Medieval & Renaissance; Baroque; Rococo; Neoclassicism; Romanticism; Realism; Impressionism & Symbolism; Fauvism & Expressionism; Cubism, Futurism & Constructivism; Dada & Surrealism; Abstract

[13] We do do not include 2007 because our data set doesn't span the full year.

Expressionism; Pop; Minimalism & Contemporary.[14] For the purposes of the current study, we create a separate subsample of transactions for each of these movements. If an artist is categorized under more than one movement, we assign all the sales of his or her work to the art movement the artist has most contributed to.[15] As a result, there is no overlap between the different subsamples, and correlations between the return estimates and volume changes across movements cannot be driven by the repeated use of the same data. The number of sales in the movement subsamples ranges from 10,485 for Neoclassicism to 102,234 for Baroque.

We make use of this dataset in two ways. First, we construct a panel data set of art returns and transaction volumes for 13 art movements. Second, we use identical resale pairs to test long-term predictions on actual transactions.

A. *Times series data*

We build aggregate and movement-specific real price indexes by applying a hedonic regression model to the full dataset and to each subsample (see Appendix VII for a description of our hedonic regression model). To construct our measure of trading volume, we record the number of observed transactions for each year in the period 1976 to 2006. Our database does not include buy-ins (i.e., items that do not reach the reserve price set by the seller), and we thus work with the numbers of lots that actually sold. We construct a proxy for the average sales rate in Section III.

Figure 3.2 presents the evolution of price and volume over our time frame. There is a strong cross-sectional correlation in prices, as previously documented by, e.g., Ginsburgh and Philippe (1995); Worthington and Higgs (2003). Interestingly, the volume series are also significantly cross-sectionally correlated. A regression of movement-level series on market-level series yields an R-squared of 0.63 for returns and 0.44 for changes in volume. Some art movements are clearly riskier than others: for example, Pop art prices culminated in 1990 to levels more than 5 times their 1984 level, *in real terms.* Our index suggests that Pop art prices subsequently fell by 83%.

[14]See Renneboog and Spaenjers (2013) for details on the compilation of the list of artists, the classification of artists into movements, and the collection of sales information.

[15]For example, Edgar Degas is classified both under Realism and under Impressionism & Symbolism. We will use the sales of his work only to estimate the price and volume changes in the market for Realism.

[Insert Figure 3.2 about here]

The riskiest art movements are shown to also be the most profitable bets. Table I presents the summary statistics on price and volume series. There is a clear risk-return relationship: the correlation between return and volatility is as high as 81.6%, as is exhibited by Panel (a) of Figure 3.3. Price and volume volatility are of the same order of magnitude; the riskier (safer) movements also exhibit more (less) volatile volume. Unsurprisingly, the riskier movements (Pop, Abstract Expressionism, Minimalism and Contemporary art) belong to the second half of twentieth century, i.e. Modern and Contemporary Art, while the safer (Romanticism, Baroque, Rococo) are Old Masters, i.e. ended no later than the first half of the nineteenth century.

We also report correlations between price and volume. The correlations are extremely high, 54% for the aggregate indices, as they are also induced by the impressive boom and bust that characterized both prices and volume during the 1990 bubble. As can be seen in Panel (b) of Figure 3.3, the riskiest schools of art are characterized by the highest correlation between prices and volumes. This pattern is reminiscent of Lee and Swaminathan (2000), who find that high (low) volume stocks exhibit many glamour (value) characteristics.[16] In our analysis, we will therefore consider subsamples using "High Volatility" (Pop, Abstract Expressionism, Minimalism and Contemporary art) and "Low Volatility" (Romanticism, Baroque, Rococo) artists.

[Insert Table I and Figure 3.3 about here]

B. Repeat-sale Data

We also use a subset of the dataset for which pairs of identical, or at least very similar objects of art can be identified. Each resale pair is considered as a unique point in our dataset, and the resales comprise 22,716 observations, spanning 1976 to 2006. For each pair of transactions, we observe the purchase and sale prices, P_i^b and P_i^s, expressed in logarithm. The log-return for holding a work of art i between the date of purchase b_i and

[16] In the absence of a significant rental market for art, we cannot compute valuation ratios such as rent to price ratios. Any definition of glamour and value is necessarily informal and based on the "test of time".

the date of sale s_i is thus given by $P_i^s - P_i^b$. The average holding period in our sample is 5.7 years.

We make use of this dataset to test the predictive relation between volume and returns, as well as price mean reversion. A potential concern is that selection bias may affect the interpretation of our results. For example, Goetzmann (1993) argues that both the upper and lower tails of art return distribution may not be observed, because works of art that fall out of fashion or are acquired by museums and major private collections are unlikely to reappear on the market. If present, such censoring is likely to be fairly small. The correlation between returns computed using a repeat-sale estimator on this subsample and the art returns using the hedonic estimator is 0.98. Both indices also show very similar long-term trends, which implies that survivorship bias is likely to be very small. Finally, the distribution of sale-to-sale returns (not shown) is quite symmetric (with a skewness of 0.27) and no particular discontinuity can be observed in the tails of the distribution.

We complete the dataset by constructing a measure of volume at the transaction level. We first collect the total number of sales on the last twelve months preceding t. Following Baker and Stein (2004), we then normalize our series by the average volume over the last five years. Taking logs, our monthly measure of volume is given by

$$\text{VOLUME}_{m,t} = \log\left(\sum_{i=t-12}^{t-1} v_{m,i}\right) - \log\left(\frac{1}{5}\sum_{j=t-60}^{t-1} v_{m,j}\right) \tag{3.1}$$

where $v_{m,t}$ is the number of transactions for movement m observed in a given month t. In order to use the largest number of observations, we assign the aggregate measure of volume to artists who are not matched to a specific movement. Aggregate volume is defined as above, but using the entire data set instead of summing the sales of a specific movement. Detrending the series brings about several benefits. First, as can be seen in Figure 3.4 for market volume, Equation (3.1) generates a persistent series. Such a property is desirable for a variable that is expected to predict long-term returns. Second, volume supposedly proxies for the price of the resale option — the overpricing component in prices — and this component must be stationary. Third, Equation (3.1) gives us a relatively high frequency series, which is not affected by art market seasonality. Finally, the series is constructed recursively, which ensures that only information that is truly available to the investor when making his forecast appears in his information set.

We merge these series with our repeat-sale dataset: for each resale pair, we record the value of $\text{VOLUME}_{m,t}$ at the month preceding the purchase and at the month preceding the sale.

[Insert Figure 3.4 about here]

Finally, we add controls for potential changes in fundamental value between the two transactions dates. As we have already emphasized the prominent role of stock market wealth effects on art prices (Hiraki et al., 2009; Goetzmann et al., 2011), we use the Global Financial Data (GFD) world index to proxy for worldwide equity wealth and equity systematic risk. In line with Mei and Moses (2005), we also include controls for other risk factors, namely the Fama-French factors Fama and French (1996) and the Pastor and Stambaugh (2003) liquidity factor. Finally, we use the one-month Treasury bill rate as the risk-free rate.

Although tastes are relatively slow-moving (Graddy, 2014), we proxy for potential changes in tastes by measuring temporal variation in artist fame. To do so, we collect the percentage of mentions of each artist name in the English-language books digitized by Google Books (Michel et al., 2011; Google, 2012). We find annual series for 2190 artists (out of 2769), which leave us with 20,604 resale pairs with taste information.

Table II gives the descriptive statistics for the repeat-sale database, expressed in log difference between the time of first and second transaction. For art, we see an average excess return of 1.2% over an average holding period of 5.7 years, with a standard deviation of 78%. Equities are undoubtedly financially dominating art, with an excess return of nearly 12% and a standard deviation of almost 29% over the period 1976-2006. Volume barely changes on average (-3%), and was much less volatile (19% standard deviation), which reflects our choice of smoothing the volume series. Finally, the percentage of mentions of each artist (fame) fell 7% on average and has a dispersed distribution (the standard deviation is 44%).

[Insert Table II about here]

Interestingly, this large volatility at the artist level averages out at movement level, as depicted by Figure 3.5. Over our sample period, we observe the increasing popularity of Andy Warhol, while the share of mentions of his name in Google Books increased by 88%. Roy Lichtenstein, another famous Pop artist, also gained increasing attention over our sample period. By contrast, Figure 3.5 shows that the exposure of the average Pop artist remained largely stable for three decades. We observe similar patterns for the other art movements: artist trajectories can be erratic, but the degree of exposure is very stable at the aggregate level. This illustrates the fact that tastes move very slowly, as pointed out by Graddy (2014), and that changes in tastes cannot explain the dramatic fluctuations that characterized art prices during that period.

[**Insert Figure 3.5 about here**]

III. The information content of volume

In a recent survey on art collecting, only a tenth of the respondents said they bought art purely as an investment, whereas 75% cited enjoyment as the key motivation (Barclays, 2012). Such a financially disinterested behavior stands in contrast with the steady growth of the art-as-an-investment industry. The specialized press regularly reports the creation of art funds, or the launch of services targeted at private investors who want to build up an art collection for investment purposes.[17] Besides surveys and anecdotal evidence, little is known about what drives volume in the fine art auction market. Since the seminal study of Baumol (1986), the literature has extensively considered what drives the demand for art, but overlooked supply and therefore volume. This makes perfect sense, as far as the supply side is thought as the *production* of works of art. Unless one discovers a forgotten masterpiece in a local flea market, the supply of works of art is inelastic — to the very least for dead artists. However, the traditional view overlooks the existence of a secondary market where collectors can sell. In a rational model such as Mandel (2009), such a secondary market doesn't exist.

In fact, the purely rational view predicts that collectors should not trade at all. This is

[17] For example, the Fine Art Fund Group, started by Philip Hoffman, a former finance director at Christie's, launched a managed art portfolio service targeted at high-net-worth individuals. The minimum investment for this type of fund is $1m over three years (Powley, 2013).

due to the "no-trade theorem": if one agent considers trading with another agent, each of them needs to consider why the other agent might be willing to trade at a particular price, which results in no trade (Tirole, 1982; Milgrom and Stokey, 1982). In order to generate trade in a rational model, one needs to assume that people are different. According to this view, people trade in the art market for idiosyncratic liquidity reasons or because their wealth or tastes differ. A change in the population of bidders can thus affect art's fundamental value and generate volume. Any increase in demand, if deemed permanent, logically leads to higher fundamental prices. Lovo and Spaenjers (2014) show that it can also lead to a contemporaneous increase in volume, even in a rational model where prices follow macroeconomic fundamentals. It is not clear, however, to what extent changes in the population of bidders can explain the sizeable price-volume correlation documented in this paper. In particular, since art yields conspicuous utility (Mandel, 2009), the utility of owning art should increase with prices, thus reducing the incentive to sell when the entry of new buyers pushes prices up.

In contrast to the purely rational view on trading, resale option models argue that speculation is an important driver of trading volume. Speculative trading arises whenever someone buys an item above its own private valuation, in order to resell it later at a higher price. A prominent example of such investment scheme is the purchase of Van Gogh's Portrait of Doctor Gachet in 1990 at the then record price of $82.5 million. The Portrait was sold within three minutes to Ryoei Saito, Japan's second-largest paper manufacturer. Immediately after taking possession of the painting, he secured it in a climate-controlled warehouse where it remained unexhibited for seven years (Taylor, 2012).

In order to distinguish consumption-driven trading from speculative trading, we look at the composition of aggregate volume. In the introduction of this paper, we showed that the share of short-term transactions peaked during the 1990 bubble. Given the huge transaction costs that characterize the art market, it is very unlikely that these works of art were bought for the pure "retinal" pleasure. These transactions thus credibly proxy for the frenzied trading predicted by Scheinkman and Xiong (2003). We construct this variable by means of the repeat-sale sample, and define it as the share of purchases that were resold within the next year. These resale pairs account for 10% to 30% of trading volume within our sample.

A common thread in the disagreement literature is that trading volume appears to act as an indicator of investor sentiment (Baker and Stein, 2004; Hong and Stein, 2007). A high volume is supposed to be symptomatic of overvaluation, signifying that the market is dominated by optimists. The following two variables therefore proxy for market sentiment.

Our first proxy for sentiment is the sales rate, the percentage of the lots sold in an auction. Sellers of individual artworks usually set a secret reserve price and if the highest bid does not reach this level, the items are "bought in" and go unsold. The convention in the art market is that the reserve price is set at or below the auctioneer's low estimate. There is anecdotal evidence that the sales rate tends to be lower in depressed markets where prices are lower and are therefore less likely to meet sellers' reserve prices (Thorncroft, 1990). Ashenfelter and Graddy (2011) find that the sales rate is not related to art prices, but is strongly positively related to unexpected price changes, defined as the difference between the hammer price and the presale estimate produced by auction house experts. A higher sales rate may therefore indicate that the market is dominated by optimists, who are willing to pay more than sellers' reserve price, which are themselves related to expert estimates. Ashenfelter and Graddy (2011) indeed report that sales rates crashed in the bust of the 1990 bubble. Since our dataset does not include items that were bought in, we construct a proxy for the sales rate. For each auction, we divide the number of observed transactions by the maximum lot number. We then take, for each year, the average sales rate across auctions as our proxy for the aggregate sales rate, from 1976 to 2006.

Newspaper articles also suggest that the share of Modern and Contemporary art is higher in "hot" markets. For example, Thorncroft (1992) reports a flight to quality (i.e. to Old Masters paintings) after the 1990 bubble burst: "the auction world has returned to its traditional ways, where connoisseurs rule and established works of art hold pride of place." This can be related to sentiment for two reasons. First, overconfident collectors are more likely to hunt for relatively young artists with larger upside potential, just as overconfident investors scrutinize the stock market hoping to find the next Google.[18] When sentiment

[18]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.'"

is low, trading should decrease and be confined to the less speculative art movements. Second, we expect collectors who engage in speculative trading to turn to the most liquid items available. Modern and contemporary items are arguably more liquid than items from older schools of art, which are to a larger extent locked up in museums or private collections.[19] A high share of modern and contemporary art in the aggregate trading volume thus signals that the market is dominated by optimistic speculators. We therefore construct a second proxy for sentiment consisting of the annual share of Modern and Contemporary art, which corresponds to our "High Volatility" group (Pop Art, Abstract Expressionism, Minimalism, and Contemporary art).

We would also like to know to what extent prices and volume correlate with changes in the population of bidders. The previous literature has emphasized the role of wealth, and in particular the wealth of the most privileged members of society, as drivers of art prices. In line with Goetzmann et al. (2011), we use the data from Piketty and Saez (2006) to build a consistent series of the share of total income received by the top 0.1 percent of all income earners in the US.

Finally, it is of independent interest to understand whether volume increases mostly because of an increase of demand, or because a higher number of items are offered for sale. Both theories predict that higher prices should be associated with a larger number of works of art offered for sale. If sentiment is high and art prices are above their fundamental values, pessimists will react and put more items for sale. Moreover, auction houses are more likely to solicit potential sellers in "hot" markets (Pesando and Shum, 2008). On the other hand, higher prices can attract sellers whose present value of future utility dividends is below the expected auction price.[20] As a proxy for the number of art objects offered for sale, we use the number of transactions divided by our sales rate.

Table III presents the correlation matrix of the percentage changes in price, volume,

[19] Modern and contemporary art should thus enjoy a higher level of liquidity, *ex ante*. A simple coordination argument suggests that they should be even more liquid *ex post*, because liquidity is self-reinforcing.

[20] A prominent example of such an increase in supply was described in the press during the 1990 bubble. At the peak of the bubble, several major museums, including the Guggenheim in New York, announced that they were disposing of important works of arts — works by Chagall, Modigliani and Kandinsky. Although this practice of de-accession is not uncommon, and serves the purpose of financing new acquisitions, it was unusual enough to be qualified a "selling spree" by the *Financial Times* (Thorncroft, 1990). The timing of the selling indeed suggests museums were trying to benefit from the very high prices reached by a few star artists (Glueck, 1990).

and the percentage changes in the five variables discussed above (the share of short-term transactions, the sales rate, the number of art objects offered for sale, the share of transactions in modern and contemporary art within the global art market, and the top income). We see that when price and volume increase, the share of short-term transactions tends to increase. The correlation between volume change and the change in short-term transactions is a highly significant 0.42.

We also learn from Table III that the sales rate tends to increase with volume. The correlation is statistically and economically large: 0.36. Volume is, by definition, given by the number of art objects offered for sale multiplied by the sales rate. If the sales rate comoves with volume, one may wonder whether volume changes *because* of changes in the sales rate. Hiraki et al. (2009) argues that the 1990 bubble was mainly driven by the influx of Japanese buyers in the art market. Anecdotal evidence also supports the idea of inexperienced and wealthy collectors competing in auction and buying at unreasonable prices. More generally, since higher bids are more likely to reach sellers' reserve prices, the entry of new buyers alone could push the sales rate up, which, in turn, would mechanically increase volume. We find, instead, that the number of art objects offered for sale increases in concert with volume. Said otherwise, people trade more, and not only because of the influx of new buyers. To see that, first consider the correlation between the sales rate and prices. Table III exhibits a 0.22 correlation, which is less than half of the price-volume correlation. Also, the former correlation is not statistically significant, in line with Ashenfelter and Graddy (2011). Moreover, if changes in volume were only due to demand shocks, the number of art objects offered for sale should not be correlated to prices. We find a highly significant 0.47 correlation between the offered art objects and prices.

[Insert Table III about here]

In line with the speculative trading hypothesis, we also expect new buyers to be primarily attracted by Modern and Contemporary art, which offer the most speculative, glamorous artists. Table III shows that the share of Modern and Contemporary art does indeed increase with both price and volume (the correlations are 0.59 and 0.52, respectively). We saw in Section II that the Modern and Contemporary Art movements are more volatile and exhibit the highest price-volume correlation. These positive correlations strengthen our argument that, when sentiment is high, buyers and sellers agree to disagree and turn to the most speculative items.

Finally, Table III provides limited support for the consumption trading hypothesis. In line with Goetzmann et al. (2011), art returns are significantly correlated with changes in inequality, with a significant 0.35 correlation, but only modestly and insignificantly with all measures of trading volume.

IV. Price-volume correlation

A. Contemporaneous relation

We now turn to the analysis of price-volume correlation across art movements, using the time series data described in Section II.A. The central relation between price and volume is presented in Table IV. Each panel reports fixed effect regression results for the 13 art movements, as well as for the 3 most volatile and 3 least volatile movements. We first look at the contemporaneous relation between prices and volumes, and then turn to regressions including equity returns, which traditionally proxy for changes in wealth.[21]

[Insert Table IV about here]

Panel A of Table IV documents a significant and pervasive price-volume correlation. Model 1 reports that a one percent change in volume is associated with an average 0.59% change in return. Price-volume correlation explains on average 25% of return variance. The relation is much stronger for high-risk movements (with an R-squared of nearly 0.40 in model 3) and much lower for Low Volatility artists (the R-squared is around 0.05, in model 5). Interestingly, the price-volume correlation remains largely intact when controlling for contemporaneous and lagged stock returns. Moreover, the R-squared increases only marginally when including stock returns, suggesting that volume is more informative than stocks in the short run.[22]

[21]Controlling for changes in top incomes instead of equity returns does not materially affect our results.

[22]Penasse (2014) presents evidence that volume is indeed more informative than stock prices for the purpose of short-term forecasting.

This result is consistent with both consumption and speculative trading, and is therefore of little help to distinguish between the theories. As argued in Section III, both hypotheses predict a causal relation from prices to the number of works of art offered for sale, and plausibly to the number of transactions. Resale option theory can also predict a lead-lag relation between lagged volume and prices, if information diffuses slowly. This is arguably the case for art where, absent a centralized market, information has to diffuse through the media and by word of mouth. Further, auctions around specific themes occur infrequently and hence the trading frequency in the art market makes information spread slowly. Empirical evidence shows that art market returns lag stock returns by a year (see e.g. Chanel (1995), Renneboog and Spaenjers (2013)), and take at least six months to reflect information contained in Sotheby's stock price (Penasse, 2014). It is therefore likely that sentiment would diffuse slowly into art prices. In the United States housing bubble, Soo (2013) similarly finds that house prices followed volume with a substantial lag and shows that both volume and prices were predicted by market sentiment.

In order to disentangle our two main hypotheses, we thus turn to an analysis of lead-lag relations between prices and volume.

B. Lead-lag relations

Before turning to the estimation of lead-lag relations, some words of caution must be spoken regarding the econometric model and the time series used. First, it is well known that fixed effects regressions with lagged dependent variables generate biased estimates. The bias is of order $1/T$ and, given the relatively large time-dimension of our panel (30 years), we do not attempt to remove it. Second, both price and volume indices are by construction moving averages over each year and may thus be artificially smooth. Time aggregation of data can lead to variances that are underestimated and autocorrelations that are overestimated relative to the true underlying process (Working, 1960). Moreover, the art auction market is very seasonal, with the second and the fourth quarters of the year generally witnessing the highest trading intensity and most important sales. In order to avoid identifying spurious lead-lag relations between price and volume, we construct distinct price and volume indices from the observations in the fourth quarter (October-December) of each year.

Panel B of Table IV shows that volume tends to lead prices. The elasticity of current returns to lagged changes in volume is 0.258 (model 1), a little less than half the elasticity to current volume, but it remains significant at the 5% level. Again, the contribution of additional predictors seems marginal: neither lagged stock returns nor lagged price remain significant when controlling for lagged volume.

Panel C tests the alternative relations where prices lead volume: we regress volume changes on lagged price changes and control in a second regression for lagged equity returns and possible volume autocorrelation. Panel C provides limited support for a causal relation from returns to volume. We only find a borderline significant relation in the High Volatility group, where a 10% price change forecasts a 2.1% change in volume on the following year (model 4). All specifications in Panel C show a negative autocorrelation of volume. A plausible explanation is market timing. Someone with a limited number of items to sell will put more items at auctions when he sees a selling opportunity, which will leave him with fewer items to dispose of the year after. A similar argument holds for a buyer with limited resources and this argument generalizes to auction houses trying to maximize revenue by soliciting more artworks when they expect higher prices.

V. Volume and overpricing

A. Volume deciles

The most important prediction of asset pricing theories incorporating disagreement is overpricing (Hong and Stein, 2007). In the absence of short selling, a high trading volume signals that prices are above the fundamental values of the art objects. A high volume at the time of a purchase should predict lower returns while controlling for potential changes in fundamental value. We test this prediction on our repeat-sale dataset, where each transaction is identified by its purchase and subsequent resale date. For each resale pair, we record the sales volume at the time of the purchase and at the time of the sale. As our analysis concentrates on overpricing, we calculate abnormal real returns, which we define as the returns in excess of the sample mean of the whole repeat-sale dataset. We construct "portfolios" based on volume deciles; the largest decile corresponds to the largest volume relative to five-year average volume. We calculate the annualized abnormal returns of resale pairs that occurred when volume was within a given decile. For example, market volume fell within the first decile from September 1982 to June 1983. We collect all pairs of transactions where the artwork was bought or sold for each decile and record the average abnormal return as a function of volume decile. Figure 3.6 exhibits the annualized abnormal returns as a function of volume; the dashed lines indicate the 5% confidence bands around the null of absence of abnormal returns.

Buying art when volume was in the highest decile yielded an average abnormal return of -3.5% per annum. This effect is economically large, compared to the average real return in our sample, which is 0.8% per annum. In contrast, a high volume of trading at the time of the sale is associated to abnormal gains of 9.3% per annum, on average. Symmetrically, low volumes at the selling date tend to be associated with low returns. For example, selling when volume is at its lowest decile generates an average abnormal loss of 8.6%. The pattern at the time of resale is stronger than the one at the time of purchase. We interpret this as evidence of overpricing, although we cannot claim that this overpricing is predictable, because the "resale" portfolios are constructed based on volume at the time of sale. Perhaps surprisingly, purchases in the lowest volume decile did not earn a significant abnormal return in the following years. The reason is that sales tended to plunge prior to the price crash when the 1990 bubble burst. Hence, a large fraction of purchases related to the first volume decile took place before prices collapsed, weighting down the average returns in that particular decile.

Ignoring the purchases in the lowest volume decile, Figure 3.6 depicts a remarkably regular pattern across deciles. Buying when volume is low and selling when volume is high seems a quite profitable strategy (but we ignore transaction costs), and is in line with the idea of bubble formation. We repeat this exercise for all subsamples and report the results in Table V. We find similar or even stronger abnormal returns for the High Volatility group, but not for the Low Volatility group, which is consistent with our previous findings on price-volume correlation.

[**Insert Figure 3.6 and Table V about here**]

B. Asset pricing models

Changes in fundamental value may be responsible for the correlations between price and volume. We therefore evaluate the impact of volume on future returns by explicitly controlling for changes in fundamental value, captured by wealth shocks and changes in tastes. In the spirit of Mei and Moses (2005), we use the classic CAPM model to estimate the systematic risk of artworks, and employ our worldwide equity index as the market index. We expand the CAPM model by our artist fame characteristic and volume. After dropping the observations from 601 artists who do not appear in Google's books database, we estimate the following equation:

$$\mathrm{r}_i - \sum_{t=b_i+1}^{s_i} r_{ft} = \alpha + \beta \sum_{t=b_i+1}^{s_i} \mathrm{MKT}_t + \gamma \sum_{t=b_i+1}^{s_i} \mathrm{FAME}_{a,t} + \nu \sum_{t=b_i+1}^{s_i} \mathrm{VOLUME}_{m,t} + \epsilon_i \quad (3.2)$$

where $\mathrm{r}_i = \sum_{t=b_i+1}^{s_i} r_{it}$ is the return on item i between b_i and s_i, computed as the difference between the log of sale price and the log of purchase price and where r_{ft} is the risk free rate. On the right hand side, we include the sum of world equity excess returns between purchase and sale times, measured by MKT_t. We also add the change in artist fame, measured by $\mathrm{FAME}_{a,t}$. We capture the change in the volume measure $\mathrm{VOLUME}_{m,t}$ as defined in Equation (3.1) for movement m. All variables are observed with monthly frequency, except $\mathrm{FAME}_{a,t}$, which is only updated annually.

Equation (3.2) states that the percentage change in the price of an artwork in excess of the risk-free rate is a function of three factors. The two fundamental factors are changes in wealth, measured by the percentage increase in the GFD equity index between the purchase and sale time, and changes in tastes measured by the increase in mentions in the Google corpus. Our test variable is ν, which measures the impact of changes in volume. If the degree of overpricing is proportional to the volume of transactions, we expect ν to be positive.

In order to control for art exposure to additional risk factors, we also extend our estimation to Fama and French (1996) factors and the Pastor and Stambaugh (2003) liquidity factor:

$$\begin{aligned}
\mathrm{r}_i - \sum_{t=b_i+1}^{s_i} r_{ft} &= \alpha + \beta \sum_{t=b_i+1}^{s_i} \mathrm{MKT}_t + \theta \sum_{t=b_i+1}^{s_i} \mathrm{SMB}_t + \phi \sum_{t=b_i+1}^{s_i} \mathrm{HML}_t \\
&+ \lambda \sum_{t=b_i+1}^{s_i} \mathrm{LIQ}_t + \gamma \sum_{t=b_i+1}^{s_i} \mathrm{FAME}_{a,t} + \nu \sum_{t=b_i+1}^{s_i} \mathrm{VOLUME}_{m,t} + \epsilon_i \quad (3.3)
\end{aligned}$$

Following, e.g., Mei and Moses (2005), we estimate Equations (3.2) and (3.3) using a three-stage estimation procedure on our sample of repeat sales, based on Case and Shiller (1987). In a first step, we regress returns on the matrix of regressors using OLS. In a second stage, we regress the squared residuals from the first step on an intercept and the time between sales. In a third step, we redo the repeat sales regression (RSR) with weighted least squares, using the fitted squared residuals as weights.

Table VI presents our empirical findings: controlling for changes in fundamental value, volume has a large positive impact on returns. The results are consistent across the samples and models and are also economically significant: for the full sample (model 1), a one-standard deviation increase in volume (19%) will increase future excess returns by 15.1% over the holding period, or by 2.6% per year. This long-term effect of volume on art returns is much larger than the effect of stock returns (that is 5.5% on average) and the effect of taste (which is 6% on average). The impact of volume is however smaller in magnitude in the subsamples, where a larger fraction of returns is captured by the market risk factor, while the models are estimated on a much smaller number of observations.

In order to ensure that these findings are not driven by a single event, namely the 1990 "bubble", we reestimate Equations (3.2) and (3.3) by allowing the effect of volume to change for each decade of our sample. For example, we estimate the CAPM as:

$$\begin{aligned}
\mathrm{r}_i - \sum_{t=b_i+1}^{s_i} r_{ft} &= \alpha\beta \sum_{t=b_i+1}^{s_i} \mathrm{MKT}_t + \gamma \sum_{t=b_i+1}^{s_i} \mathrm{FAME}_{a,t} \\
&+ \nu_1 \sum_{t=b_i+1}^{1986} \mathrm{VOLUME}_{m,t} + \nu_2 \sum_{1987}^{1996} \mathrm{VOLUME}_{m,t} + \nu_3 \sum_{1997}^{s_i} \mathrm{VOLUME}_{m,t} + \epsilon_i
\end{aligned}$$

Table VII presents the estimated values for each decade: the coefficients associated with volume are again significant in each specification. Perhaps unsurprisingly, the impact of volume is much larger during the 1987-1996 period, which coincides with the Japanese bubble. In the last decade of our sample period, the impact of volume is less than half the coefficient estimated on the full sample, but we must keep in mind that we only observe transactions that occurred over a limited time span (purchases that took place in the beginning of that decade given that the holding period averages more than 6 years). Moreover, prices soared after 2001 and our dataset do not include transactions from the 2008-2009 price collapse.

[**Insert Tables VI and VII about here**]

VI. Discussion

This paper documents evidence supporting theories where agents engage in speculative trading that pushes art price above fundamentals, which are defined, for a given work of art, as the private valuation of the most optimistic agent. A key feature of these models is that pessimists cannot sell short, which implies that their own valuations are not incorporated into prices. Our main findings can be summarized as follows: (i) a high trading volume coincides with more speculative trades and higher market sentiment, (ii) prices and volume are significantly correlated, and this correlation is higher in the most volatile segments of the art market, and (iii) a high volume predicts negative returns. We readily argued that these findings are difficult to reconcile with a model where agents hold identical beliefs and trade on the basis of taste or wealth. In this section, we examine alternative mechanisms that can generate some of these findings and discuss alternative interpretations.

A. Time-varying risk premia

The most straightforward explanation of return predictability is that risk premia or discount rates vary over time (Cochrane, 2011). To the extent that our volume measure correlates with discount rates, the predictability we observe is consistent with a rational model. For example, volume may vary with business cycles fluctuations, such that when volume is low, art collectors demand higher risk premia for holding works of art. It seems however fairly unlikely that art volume would capture business cycle fluctuations or risk premia, which are not readily captured by the four risk factors we include in our pricing model. Moreover, we readily argued that such a model would be unable to explain the composition of volume, for example that the share of short-term transactions is significantly related to volume and prices. Finally, we showed that when volume is very high, collectors on average earned negative returns, which is incompatible with the assumption that predictability reflects a risk premium.

B. *Alternative bubble models*

In a seminal article, Blanchard and Watson (1983) provide a bubble model that is fully consistent with rational expectations and constant expected returns. The overpricing component in asset prices is independent of the asset's fundamental and bursts on any period with a constant probability. If the bubble does not burst, it grows at a faster rate than the discount rate. It may therefore be rational to ride a bubble, if it grows on average at the same rate as the discount rate. Although rational bubbles can occur in infinite horizon models, the theoretical conditions needed to support them are quite stringent and generally require that the asset price bubble not emerge over time. If, for instance, new works of art are created when prices increase, and if these new works are viewed as appropriate substitutes for existing works of art, no bubble can emerge.[23] An important difference between rational bubble models and the resale option theory is that rational bubbles must grow explosively, while the bubble component generated by the resale option is stationary over time. It is this stationary component that fuels return predictability, while rational bubble models generally assume that expected returns are constant over time.

An important strand of the behavioral finance literature suggests that investors are likely to form expectations by extrapolating past price changes, which may generate bubbles. This idea features prominently in many classical accounts of asset bubbles (e.g. Kindleberger (1978), Minsky (1986), Shiller (2000)). Investors extrapolate past prices because they suffer from behavioral biases such as the representativeness bias, a tendency to

[23]See e.g. Brunnermeier and Oehmke (2013) for further discussion on the theoretical conditions required to support rational bubbles.

view events as typical of some specific class (Barberis et al., 1998), or the self-attribution bias, a tendency to attribute success to one's own ability but failure to external factors (Daniel et al., 1998). A central prediction of these models is time series momentum: returns should be positively autocorrelated. Table IV provides little support for this prediction, where returns are not found to be significantly autocorrelated when controlling for lagged volume. Furthermore, returns are only positively autocorrelated for the High Volatility art movements (model 4).

C. Credit and leverage cycles

A frequent explanation of asset prices bubbles emphasizes the role of credit and leverage. This is particularly true for the real estate market, most people borrow to buy houses. Stein (1995) argues that, because the ability to borrow is directly tied to the value of houses, a positive income shock that increases housing demand and real estate prices relaxes the borrowing constraint, which further increases the demand for houses. In a heterogeneous-beliefs model, Geanakoplos (2010) introduces a mechanism where pessimistic agents are willing to finance investments made by optimistic agents. In contrast to resale option models, beliefs do not change over time, and therefore there is no resell premium. Large price fluctuations are attributed to fluctuations in endogenous leverage, which is too high in booming periods and too low in declining periods.

These credit and leverage models predict that changes in wealth may have more-than-proportional effects on asset prices, while resale option models emphasize the role of fluctuations in beliefs. A distinct feature of the art market is that credit and leverage play almost no role. Auction houses may sometimes lend part of the purchase to buyers, but this practice is uncommon and confined to major purchases (Thompson, 2009).[24] It is therefore quite unlikely that art price fluctuations are driven by credit cycles.

[24] Auction houses can also provide guarantees to sellers who are concerned that not enough bidders will enter the auctions for their items. Such guarantees can also be provided by third parties. Graddy and Hamilton (2014) study the effect of guarantees (both in-house and third party) and find that they have no economic effect on final prices.

D. Loss aversion

The disposition effect (Shefrin and Statman, 1985) can also explain a positive price-volume relation. If loss averse collectors decide to delay losses and choose not to sell in falling markets, then price declines could predict a decrease in trading volume. In a framework close to ours, Clayton et al. (2008) provide strong support for this hypothesis in the US real estate market, where real-estate prices are shown to Granger-cause trading volume. Statman et al. (2006) argue that Granger causality of prices towards volume can also be generated by overconfidence about one's trading skills and they provide evidence for both the disposition effect and overconfidence in the stock market.

If loss aversion is driving the price-volume correlation in the art market, we expect lagged returns to have some forecasting power and negative returns to predict drops in volume. This is precisely what we find in Table VIII, which reports that art market losses forecast low volumes. For the full sample (model 1), a 10% drop in art prices is related to an average decline of 1.8% in volume, an effect significant at the 5% level. We find that this relation is strongest for the most risky art movements (model 2).

These findings are in line with Graddy et al. (2014) who provide evidence of loss aversion in repeat-sale art data. The disposition effect is therefore likely to amplify the fall in volume following market crashes, but is by definition unable to drive the contemporaneous booms in both prices and volumes.

[Insert Table VIII about here]

E. Volume and the cross section of art returns

It is important to emphasize that our results pertain to the time series dimension of art returns. We find little evidence of a cross-sectional relation between trading volume and returns. Although some art movements are more speculative than others, we do not find that, within a given year, high volume artists tend to earn lower returns. Fads can temporarily affect the prices of some artists, as shown by Penasse et al. (2014), but they are not systematically characterized by large volumes. To put it differently, our findings suggest that changes in volume reflect shocks that affect the market *as a whole*, where "fads" pertain to the cross section of art returns.

To illustrate this, we use the repeat-sale dataset and construct three portfolios based on trading volume. For each year of a purchase, we sort transactions based on the volume variable defined in section II and we allocate to the first portfolio all purchases falling in the lowest volume tercile, to the second portfolio the transactions from the second tercile and to the third portfolio the transactions with the highest volume. In each case, we apply the RSR methodology to estimate returns; in other words, we "buy" and "sell" whenever the owner bought and sold in reality. We show the evolution of each price index in Figure 3.7. If high volume artists where overpriced with respect to low volume artists, we would expect the "high volume" strategy to underperform in the long run.[25] Although the "high volume" strategy appears to be more volatile than the two others, especially during the peak of 1990, we find no evidence of underperformance. For example, the p-value of a t-test on the difference between "high volume" and "low volume" returns is 0.14. In contrast to the results of Section V, where we allocated sales based on the time series dimension, the strategies based on the cross-sectional dimension show little heterogeneity. This suggests that volume contains little information about whether a given art movement is subject to a temporary fad. Indeed, we readily argued that volume across art movements is mostly affected by common shocks, which we attributed to changes in market sentiment. This result also rules out the interpretation of the predictive results of Section V in terms of liquidity premia. If art collectors were willing to pay a premium for the most liquid (and volatile) art movements, we would plausibility obtain a cross-sectional relation between trading volume and returns.

[Insert Figure 3.7 about here]

VII. Conclusion

The art market is subject to large fluctuations that characterize prices and trading volume, and that are difficult to reconcile with a rational model which captures how people trade to consume. This paper argues that limits to arbitrage, namely the impossibility to sell

[25]The underperformance of high-volume stocks is pervasive in the literature (see e.g. Brennan et al. (1998), Datar et al. (1998)) and is generally interpreted as evidence for a liquidity premium and differences of opinion.

art short, induce a speculative component to art prices. As pessimists cannot short-sell, their opinions are not incorporated into art prices, which hence only reflect the opinion of the most optimistic collectors. As a result, an optimist is willing to pay more than her own private value because she knows that, in the future, there may be other collectors that value the work of art more than she does. The difference between her willingness to pay and her own private value reflects a speculative motive, the value of the right to sell the work of art in the future.

This paper investigates this theory by looking at the behavior of art prices and volumes and by directly measuring returns over a comprehensive data set of worldwide art auctions. Rising prices tend to be accompanied by more short-term transactions, which we interpret as trading frenzies, given the huge trading costs that characterize the art market. Trading frenzies tend to concentrate on the works of Modern and Contemporary artists, for which prices and volume are on average more volatile and more correlated. When trading volume is high, we find that buyers tend to overpay, in that a high volume strongly predicts negative returns in the subsequent years. Art returns are therefore predictable, not because risk premia change over time as in traditional models, but instead because prices fluctuate above the fundamental value that would prevail in the absence of short-selling constraints.

Appendix: Hedonic regression

Hedonic regressions are a popular methodology for constructing constant-quality price indexes for infrequently traded goods like houses or collectibles. Hedonic models control for temporal variation in the quality of the transacted goods by attributing implicit prices to their "utility-bearing characteristics" (Rosen, 1974). Our model relates the natural logs of USD hammer prices to quarterly dummies, while controlling for a wide range of hedonic characteristics. More formally, our regression can be expressed as follows:

$$\ln P_{kt} = \alpha + \sum_{m=1}^{M} \beta_m X_{mkt} + \sum_{t=1}^{T} \gamma_t D_{kt} + \epsilon_{kt} \tag{3.4}$$

where P_{kt} represents the real USD price of an art object k at time t, X_{mkt} is the value of characteristic m of item k at time t, and D_{kt} is a dummy variable that equals one if object k is sold in time period t. The coefficients β_m reflect the attribution of a relative shadow price to each of the m characteristics. The estimates of γ_t can be used to construct an art price index.[26] Apart from the variables related to the timing of the sale, the hedonic variables X used are the same as in Renneboog and Spaenjers (2013), and capture characteristics of the artist (through the inclusion of artist dummies and an art history textbook dummy), the work (through the inclusion of variables capturing attribution, authenticity, medium, size, and topic), and the sale (through the inclusion of auction house dummies). The R-squareds of the different hedonic regressions lie between 56% and 76% (detailed results available on request).

[26] A subtle point is that the resulting index tracks the geometric means — not the arithmetic means — means of prices over time, due to the log transformation prior to estimation.

Bibliography

Aït-Sahalia, Yacine, Jonathan A. Parker, and Motohiro Yogo, 2004, Luxury goods and the equity premium, *Journal of Finance* 59, 2959–3004.

Ashenfelter, Orley, and Kathryn Graddy, 2011, Sale rates and price movements in art auctions, *American Economic Review* 101, 212–216.

Bai, Jennie, Jia Guo, and Benjamin R. Mandel, 2013, *Going global: markups and product quality in the Chinese art market*, Staff report no. 614, Federal Reserve Bank of New York.

Baker, Malcolm, and Jeremy C. Stein, 2004, Market liquidity as a sentiment indicator, *Journal of Financial Markets* 7, 271–299.

Barberis, Nicholas, Andrei Shleifer, and Robert Vishny, 1998, A model of investor sentiment, *Journal of Financial Economics* 49, 307–343.

Barclays, 2012, Profit or pleasure? Exploring the motivations behind treasure trends, *Barclays Wealth Insights* 15.

Baumol, William J., 1986, Unnatural value: or art investment as floating crap game, *American Economic Review* 76, 10–14.

Beggs, Alan, and Kathryn Graddy, 2009, Anchoring effects: evidence from art auctions, *American Economic Review* 99, 1027–1039.

Biais, Bruno, and Peter Bossaerts, 1998, Asset prices and trading Volume in a beauty contest, *Review of Economic Studies* 65, 307–340.

Blanchard, Olivier J., and Mark W. Watson, 1983, *Bubbles, rational expectations and financial markets*, Working paper 945, National Bureau of Economic Research, Cambridge, MA.

Brennan, Michael J., Tarun Chordia, and Avanidhar Subrahmanyam, 1998, Alternative factor specifications, security characteristics, and the cross-section of expected stock returns, *Journal of Financial Economics* 49, 345–373.

Brunnermeier, Markus K., and Martin Oehmke, 2013, Bubbles, financial crises, and systemic risk, in George M. Constantinides, Milton Harris, and Rene M. Stulz, eds., *Handbook of the Economics of Finance*, volume 2, chapter 18, 1221–1288 (Elsevier).

Carlos, Ann, Larry Neal, and Kirsten Wandschneider, 2006, *Dissecting the anatomy of exchange alley: the dealings of stockjobbers during and after the South Sea bubble*, Unpublished paper, University of Illinois.

Case, Karl E., and Robert J. Shiller, 1987, Prices of single-family homes since 1970: New indexes for four cities, *New England Economic Review* 45–56.

Chanel, Olivier, 1995, Is art market behaviour predictable?, *European Economic Review* 39, 519–527.

Clayton, Jim, Norman Miller, and Liang Peng, 2008, Price-volume correlation in the housing market: causality and co-movements, *Journal of Real Estate Finance and Economics* 40, 14–40.

Cochrane, John H., 2003, Stocks as money: convenience yield and the tech-stock bubble, in William C. Hunter, George G. Kaufman, and Michael Pomerleano, eds., *Asset Price Bubbles: The Implications for Monetary, Regulatory, and International Policies* (MIT Press, Cambridge, MA).

Cochrane, John H., 2011, Presidential address: discount rates, *Journal of Finance* 66, 1047–1108.

Cujean, Julien, and Michael Hasler, 2014, *Disagreement and predictability over the business cycle*, Working paper, University of Maryland.

Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam, 1998, Investor psychology and security market under- and overreactions, *Journal of Finance* 53, 1839–1885.

Datar, Vinay T., Narayan Y. Naik, and Robert Radcliffe, 1998, Liquidity and stock returns: An alternative test, *Journal of Financial Markets* 1, 203–219.

De Silva, Dakshina G., Rachel A.J. Pownall, and Leonard Wolk, 2012, Does the sun shine on art prices?, *Journal of Economic Behavior & Organization* 82, 167–178.

Duthy, Robin, Art as investment - strength in depth, 1989. Financial Times.

Fama, Eugene F., and Kenneth R. French, 1996, Multifactor explanations of asset pricing anomalies, *Journal of Finance* 51, 55–84.

Favara, Giovanni, and Zheng Song, 2013, House price dynamics with dispersed information, *Journal of Economic Theory* 1, 1–33.

Geanakoplos, John, 2010, The Leverage Cycle, *Cowles Foundation Discussion Paper* 1304.

Genesove, David, and Christopher J. Mayer, 2001, Loss aversion and seller behavior: evidence from the housing market, *Quarterly Journal of Economics* 116, 1233–1260.

Ginsburgh, Victor, and Jeanfils Philippe, 1995, Long-term comovements in international markets for paintings, *European Economic Review* 39, 538–548.

Ginsburgh, Victor, and Sheila Weyers, 2008, On the contemporaneousness of Roger de Piles' Balance des Peintres, in Stephen E. Cullenberg, Joseph W. Childers, and Jack Amariglio, eds., *Sublime Economy: On the Intersection of Art and Economics*, volume 2008, chapter 5, 112–123 (Routledge, London).

Glueck, Grace, The mania of art auctions: problems as well as profits, 1988. The New York Times.

Glueck, Grace, Guggenheim to sell 3 works to help buy others, 1990. The New York Times.

Goetzmann, William N., 1993, Accounting for taste: Art and the financial markets over three centuries, *American Economic Review* 83, 1370–1376.

Goetzmann, William N., Luc Renneboog, and Christophe Spaenjers, 2011, Art and money, *American Economic Review* 101, 222–226.

Google, Google Books Ngram Viewer, 2012.

Graddy, Kathryn, 2014, Taste Endures! The rankings of Roger de Piles (1709) and three centuries of art prices, *Journal of Economic History* 73, 766–791.

Graddy, Kathryn, and Jonathan Hamilton, 2014, *Auction house guarantees for works of art*, Discussion paper no. 9996, Centre for Economic Policy Research.

Graddy, Kathryn, Lara Loewenstein, Jianping Mei, and Mike Moses, 2014, Anchoring or loss aversion? Empirical evidence from art auctions, *Association for Cultural Economics International Working Paper* AWP–04–201.

Harris, Milton, and Artur Raviv, 1993, Differences of opinion make a horse race, *Review of Financial Studies* 6, 473.

Harrison, J. Michael, and David M. Kreps, 1978, Speculative investor behavior in a stock market with heterogeneous expectations, *Quarterly Journal of Economics* 92, 323.

Hiraki, Takato, Akitoshi Ito, Darius A. Spieth, and Naoya Takezawa, 2009, How did Japanese investments influence international art prices?, *Journal of Financial and Quantitative Analysis* 44, 1489.

Hong, Harrison, and Jeremy C. Stein, 2007, Disagreement and the stock market, *Journal of Economic Perspectives* 21, 109 –128.

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

Korteweg, Arthur, Roman Kräussl, and Patrick Verwijmeren, 2013, Does it pay to invest in art? A selection-corrected returns perspective, *LSF Research Working Paper Series* 13–07.

Kräussl, Roman, Thorsten Lehnert, and Nicolas Martelin, 2014, *Is there a bubble in the art market?*, Working paper, Luxembourg School of Finance, University of Luxembourg.

Lee, Charles M. C., and Bhaskaran Swaminathan, 2000, Price momentum and trading volume, *Journal of Finance* 55, 2017–2069.

LeRoy, Stephen F., 2004, Rational exuberance, *Journal of Economic literature* XLII, 783–804.

Lovo, Stefano, and Christophe Spaenjers, 2014, *Unique durable assets*, Research paper no. fin-2014-1037, HEC Paris.

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

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

Mei, Jianping, Jose A. Scheinkman, and Wei Xiong, 2009, Speculative trading and stock prices: Evidence from Chinese AB share premia, *Annals of Economics and Finance* 10, 225–255.

Michel, Jean-Baptiste, Yuan Kui Shen, Aviva Presser Aiden, Adrian Veres, Matthew K. Gray, Joseph P. Pickett, Dale Hoiberg, Dan Clancy, Peter Norvig, Jon Orwant, Steven Pinker, Martin A. Nowak, and Erez Lieberman Aiden, 2011, Quantitative analysis of culture using millions of digitized books., *Science* 331, 176–182.

Milgrom, Paul R., and Nancy Stokey, 1982, Information, trade and common knowledge, *Journal of Economic Theory* 12, 112–128.

Miller, Edward M., 1977, Risk, uncertainty, and divergence of opinion, *Journal of Finance* 50, 1013.

Minsky, Hyman P., 1986, *Stabilizing an unstable economy* (Yale University Press, New Haven).

Ofek, Eli, and Matthew Richardson, 2003, DotCom mania: the rise and fall of internet stock prices, *Journal of Finance* 58, 1113–1138.

Palfrey, Thomas R., and Stephanie W. Wang, 2012, Speculative overpricing in asset markets with information flows, *Econometrica* 80, 1937–1976.

Pastor, Lubos, and Robert F. Stambaugh, 2003, Liquidity risk and expected stock returns, *Journal of Political Economy* 111, 642–685.

Pastor, Lubos, and Pietro Veronesi, 2006, Was there a Nasdaq bubble in the late 1990s?, *Journal of Financial Economics* 81, 61–100.

Peers, Alexandra, Will economy push art world off the ledge? — Auctions will test resilience of boom; prizing the prolific, 2008. The Wall Street Journal Europe.

Penasse, Julien, 2014, *Real-time forecasts of auction prices*, Working paper, Tilburg University.

Penasse, Julien, Luc Renneboog, and Christophe Spaenjers, 2014, Sentiment and art prices, *Economics Letters* 122, 432–434.

Pesando, James E., and Pauline M. Shum, 2007, The law of one price, noise and irrational exuberance: the auction market for Picasso prints, *Journal of Cultural Economics* 31, 263–277.

Pesando, James E., and Pauline M. Shum, 2008, The auction market for modern prints: Confirmations, contradictions, and new puzzles, *Economic Inquiry* 46, 149–159.

Piazzesi, Monika, and Martin Schneider, 2009, Momentum traders in the housing market: survey evidence and a search model, *American Economic Review* 99, 406–411.

Piketty, Thomas, and Emmanuel Saez, 2006, *Income inequality in the United States, 1913-2002*, Working paper, University of California, Berkeley.

Pouget, Sebastien, Julien Sauvagnat, and Stephane Villeneuve, 2014, A mind is a terrible thing to change: confirmatory bias in financial markets, *IDEI Working Paper Series* 720.

Powley, Tanya, Is it time to put money into Manet or Monet?, 2013. The Financial Times.

Pownall, Rachel A.J., Stephen Satchell, and Nandini Srivastava, 2013, *A random walk through Mayfair : dynamic models of UK art market prices and their dependence on UK equity prices*, Working paper, Maastricht University.

Renneboog, Luc, and Christophe Spaenjers, 2013, Buying beauty: on prices and returns in the art market, *Management Science* 59, 36–53.

Rosen, Sherwin, 1974, Hedonic prices and implicit markets: product differentiation in pure competition, *Journal of Political Economy* 82, 34–55.

Scheinkman, Jose A., and Wei Xiong, 2003, Overconfidence and speculative bubbles, *Journal of Political Economy* 111, 1183–1220.

Shefrin, Hersh, and Meir Statman, 1985, The disposition to sell winners too early and ride losers too long: Theory and evidence, *Journal of Finance* 777–790.

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

Soo, Cindy K., 2013, *Quantifying animal spirits: news media and sentiment in the housing market*, Working paper, University of Michigan.

Statman, Meir, Steven Thorley, and Keith Vorkink, 2006, Investor overconfidence and trading volume, *Review of Financial Studies* 19, 1531–1565.

Stein, Jeremy C., 1995, Prices and trading volume in the housing market: a model with down-payment effects, *Quarterly Journal of Economics* 110, 379–406.

Stigler, George J., and Gary S. Becker, 1977, De gustibus non est disputandum, *American Economic Review* 67, 76–90.

Taylor, Mark C., Is modern finance ruining modern art?, 2012. Bloomberg.

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

Thorncroft, Antony, Arts: Selling spree in US, 1990. The Financial Times.

Thorncroft, Antony, Old Masters lead the way to normality - A year in which tradition - and sanity - returned to the auction houses, 1992. The Financial Times.

Tirole, Jean, 1982, On the possibility of speculation under rational expectations, *Econometrica: Journal of the Econometric Society* 50, 1163–1181.

Vogel, Carol, Warhols of tomorrow are dealers' quarry today, 2006. The New York Times.

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

Worthington, Andrew C., and Helen Higgs, 2003, Art as an investment: short and long-term comovements in major painting markets, *Empirical Economics* 649–668.

Xiong, Wei, 2013, Bubbles, crises, and heterogeneous beliefs, in J. P. Fouque, and J. A. Langsam, eds., *Handbook for Systemic Risk*, chapter 24, 663–713 (Cambridge University Press).

Xiong, Wei, and Jialin Yu, 2011, The Chinese warrants bubble, *American Economic Review* 101, 2723–2753.

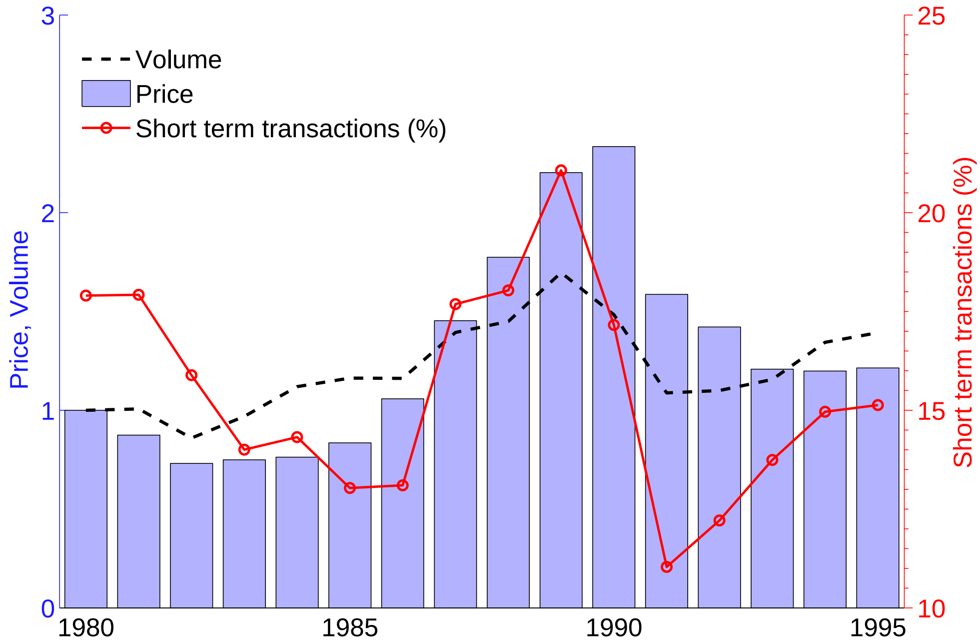

Figure 3.1: Price, Volume and Short Term Transactions During the 1990 Bubble
This figure shows aggregate art prices, the total volume of transactions, and the share of short-term transactions during the 1990 bubble (1980-1995). Prices and volume are expressed in function of their 1980 level (left scale). The share of short-term transactions is defined as the share of purchases that were resold within the next year, and is computed from the repeat-sale data set (right scale).

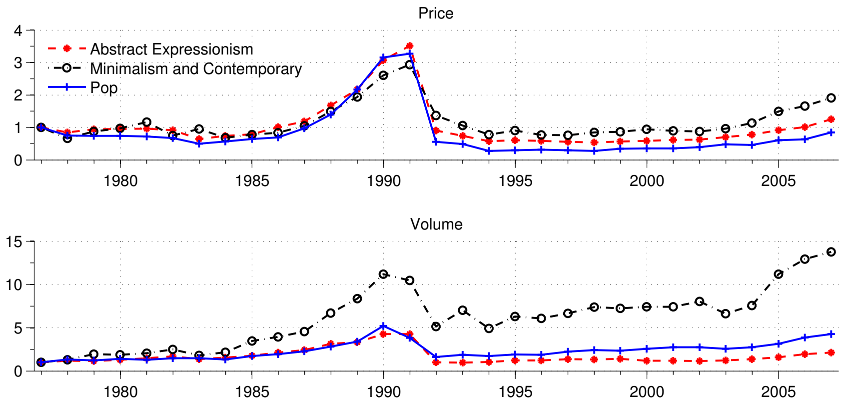

(a) High Volatility movements

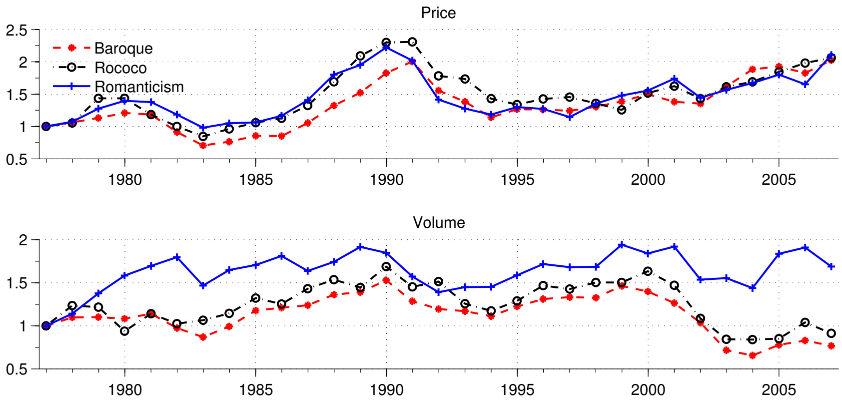

(b) Low Volatility movements

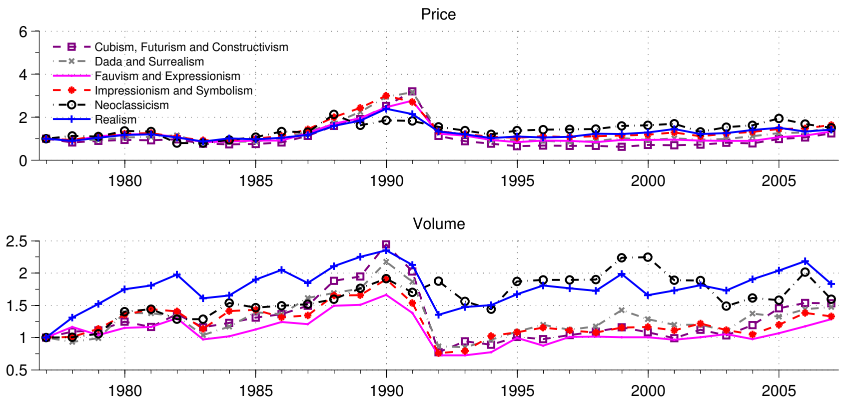

(c) Remaining movements

Figure 3.2: Prices and Volumes of 13 Art Movements (1976-2006)
This figure plots prices and volume for each of the thirteen art movements, with 1976 as their standardized benchmark level.

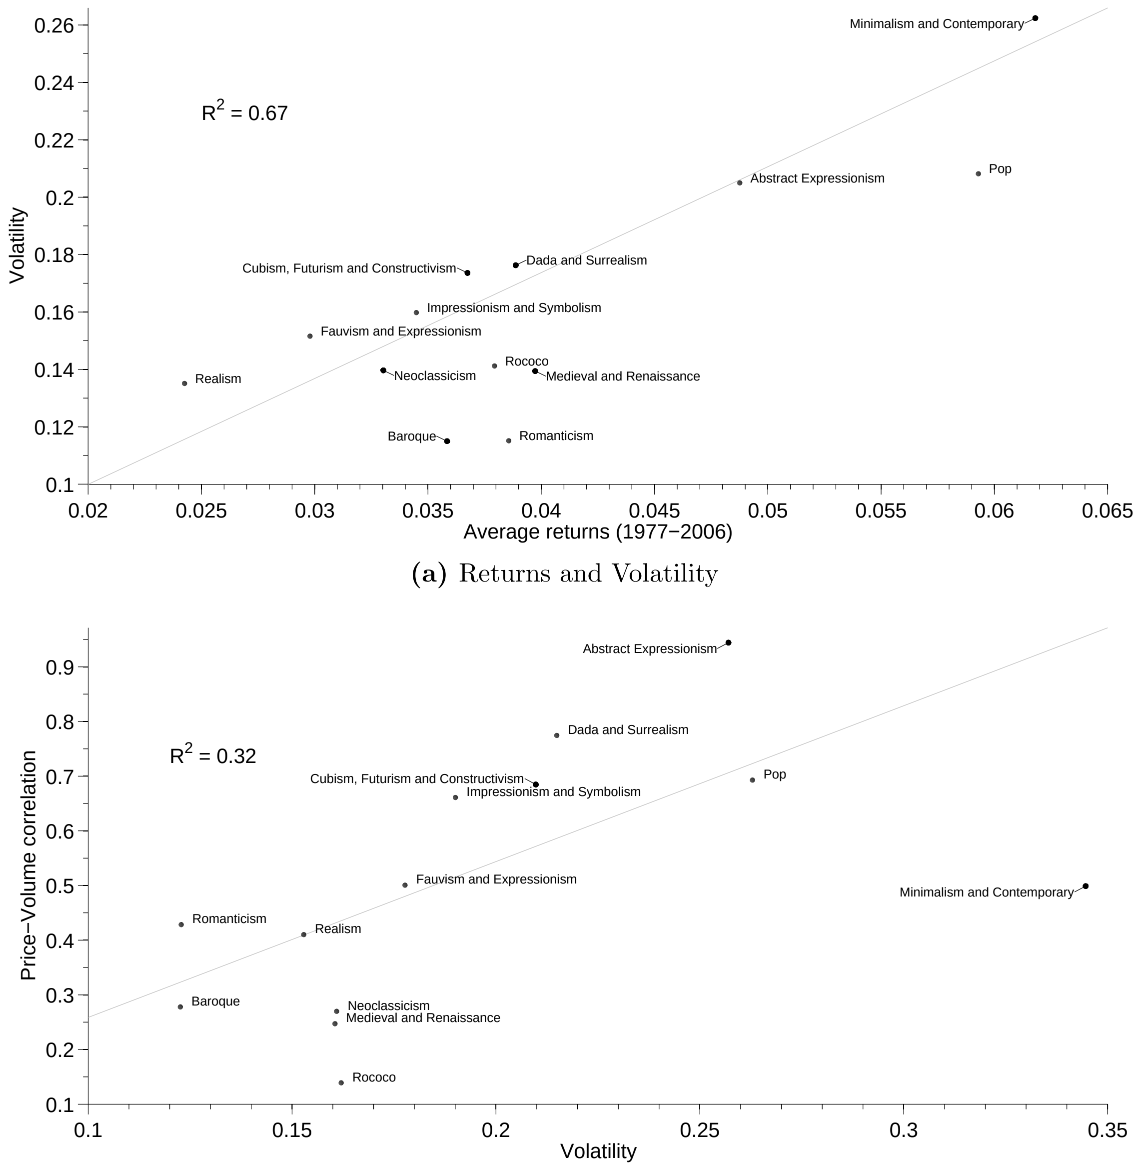

(a) Returns and Volatility

(b) Volatility and Price-Volume Correlation

Figure 3.3: Returns, Volatility and Price-Volume Correlation
These figures are scatter-plots of the first and second moments of the return series (the values of which are provided in Table I). Panel (a) plots the average return of each of the thirteen art movements of our sample against their volatility. Panel (b) plots the volatility of each movement against their price-volume correlation.

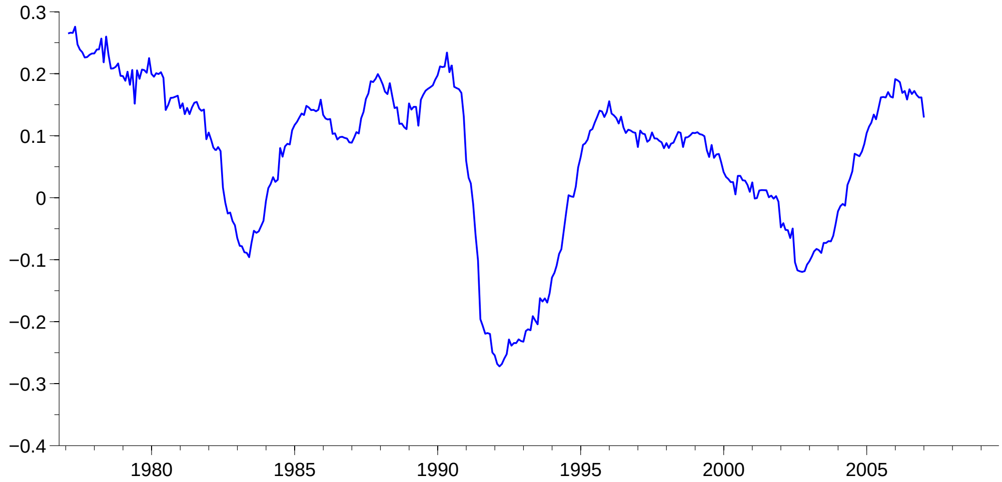

Figure 3.4: Detrended market volume
This figure plots the monthly measure of volume constructed by means of Equation (3.1) and the whole sample of 1.1 million of auction transactions. Each month t we take the log of the total number of sales on the last twelve months preceding t. We then normalize our series by subtracting the log of the average number of sales over the last five years.

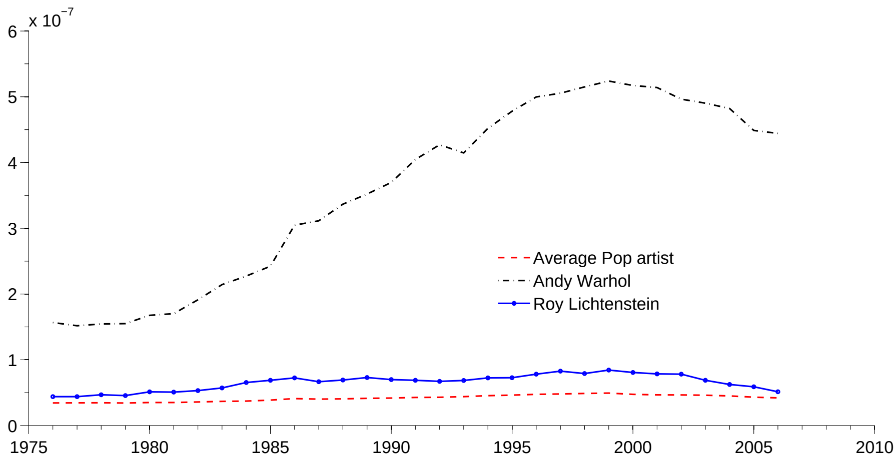

Figure 3.5: Artist Fame: Andy Warhol and Pop Art
This figure depicts the share of mentions of Andy Warhol's and Roy Lichtenstein's names in the Google Books database, and of the average share of the 111 Pop artists in our data set.

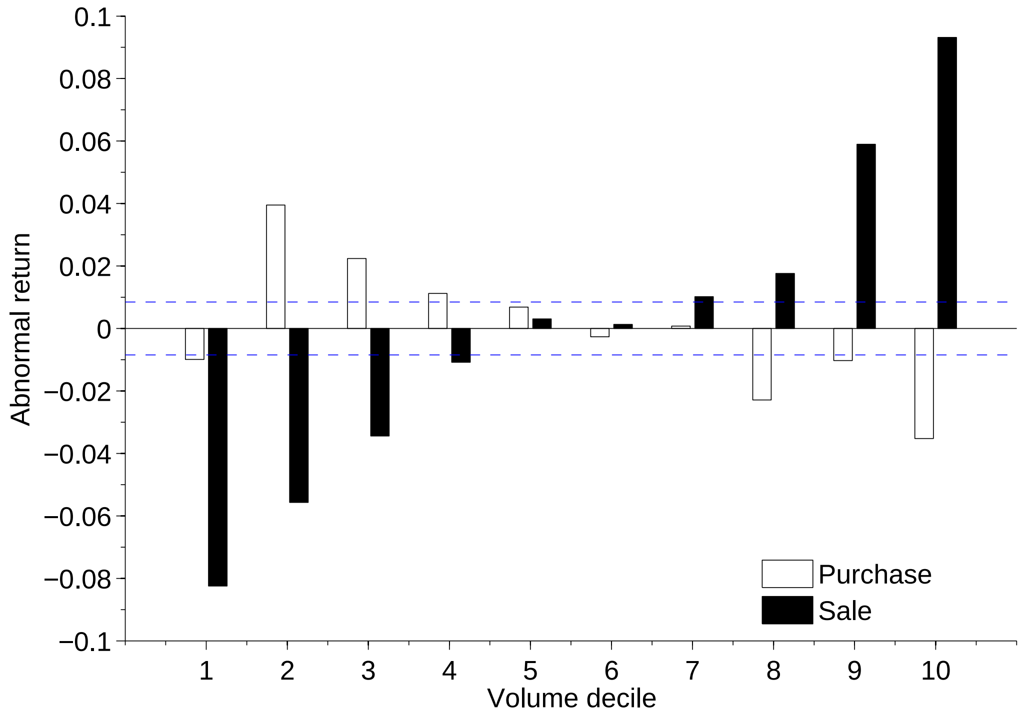

Figure 3.6: Volume and Abnormal Returns
This figure shows abnormal returns expressed by volume deciles (whereby decile 10 corresponds to the largest volume). Each repeat-sale transaction is identified by purchase and sale date. We construct "portfolios" of paintings based on volume at the time of purchase or sale. Volume is constructed according to Equation (3.1) (see also Figure 3.4). We compute the average abnormal return of each "portfolio" and rank them from the lowest volume (first decile) to highest (tenth decile) for the full sample and specific subgroups. Abnormal returns are defined as annualized returns in excess of the sample average. The dashed lines indicate the 5% confidence bands around the null of absence of abnormal returns.

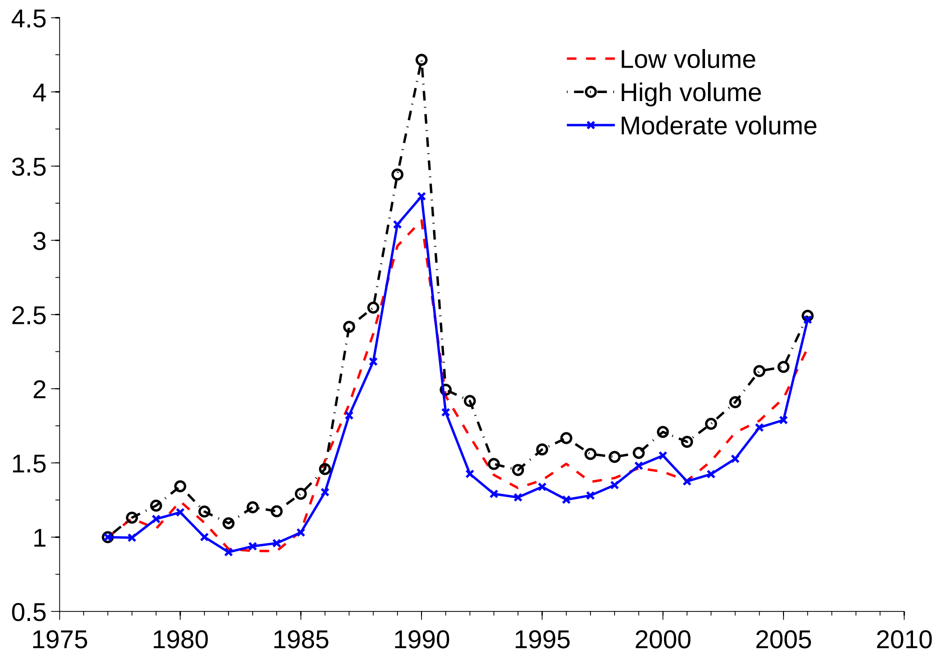

Figure 3.7: Repeat-Sales Price Indices
This figure presents the prices indices constructed from the "Low Volume", "Moderate Volume" and "High Volume" strategies (described in Section VI.E). The indices are obtained by applying a repeat-sale regression to resale pairs allocated to each volume tercile.

Table I: Time Series Data: Sample statistics

	Price changes				Volume changes				Price-volume correlation
	Mean	S.D.	Min	Max	Mean	S.D.	Min	Max	
Abstract Expressionism	4.2	22.3	-74.1	41.7	5.7	19.5	-76.4	30.8	78.0
Baroque	3.2	13.2	-22.9	26.0	-0.2	11.3	-30.7	18.8	24.1
Cubism, Futurism, and Constructivism	3.3	20.7	-64.9	40.6	3.3	16.6	-61.7	26.0	56.5
Dada and Surrealism	3.5	18.7	-55.8	30.8	2.9	16.9	-53.7	39.1	64.0
Fauvism and Expressionism	2.8	17.9	-55.6	38.5	2.1	14.6	-47.6	28.2	42.0
Impressionism and Symbolism	3.1	16.4	-50.4	39.5	2.3	15.4	-50.6	29.3	54.9
Minimalism and Contemporary	5.3	23.7	-53.4	42.6	12.1	24.7	-50.9	59.8	41.8
Medieval and Renaissance	3.5	17.7	-34.8	50.3	0.4	13.6	-38.5	27.9	21.6
Neoclassicism	3.0	18.8	-40.0	62.7	2.4	13.6	-21.0	32.2	23.5
Pop	5.0	27.7	-83.1	55.9	7.2	19.7	-58.1	53.2	57.4
Realism	2.3	15.2	-38.2	35.6	2.9	13.1	-36.4	30.9	34.7
Rococo	3.4	14.2	-22.8	36.0	0.6	13.7	-26.1	23.5	13.0
Romanticism	3.4	13.7	-29.8	28.2	2.4	11.3	-19.9	27.7	36.2
Art market	3.6	13.8	-32.0	37.4	3.8	10.3	-26.7	20.0	54.3

This table presents the descriptive statistics (mean, standard deviation (S.D.), minimum, maximum, and price-volume correlation) of the log-differences of prices and volumes, for each of the thirteen movements and for the whole art market, expressed in percentage terms.

Table II: Repeat-sale Data: Sample Statistics

	Mean	S.D.	Min	Max
ART	1.20	78.43	−458.20	554.05
EQ MKT	11.97	28.85	−77.86	128.06
SMB	4.87	21.64	−61.43	89.12
HML	24.12	29.60	−57.08	164.11
LIQ	31.13	38.55	−28.20	198.95
FAME	−6.03	40.49	−367.89	401.86
VOLUME	−2.80	19.46	−107.82	74.91
N	24889			

This table presents the descriptive statistics (mean, standard deviation (S.D.), minimum and maximum) of the variables used in our repeat-sale analysis. All variables are expressed in percentage changes between each resale pair. ART is the return on artworks between the purchase and sale times in excess of the risk-free rate: $\text{ART}_i = \sum_{t=b_i+1}^{s_i} r_{it} - \sum_{t=b_i+1}^{s_i} r_{ft}$. EQ MKT measures equity excess returns, and SML, HML and LIQ are the Fama and French (1996) and Pastor and Stambaugh (2003) risk factors. FAME is the share of mentions in Google Books for each artist and VOLUME is the volume measure, defined in Equation (3.1).

Table III: Information Content of Trading Volume

	Price	Volume	Short term trans.	Sales rate	Art obj. off. for sale	Mod. and Contemp.	Top inc.
Price	1.00						
Volume	**0.54**	1.00					
Short-term transactions	**0.35**	**0.58**	1.00				
Sales rate	0.22	**0.36**	0.19	1.00			
Art objects offered for sale	**0.47**	**0.86**	**0.52**	-0.15	1.00		
Modern and Contemporary	**0.59**	**0.52**	**0.50**	0.02	**0.57**	1.00	
Top income	**0.35**	0.23	0.29	0.22	0.14	0.26	1.00

This table presents pairwise correlations between price, volume, the share of short-term transactions, the sales rate, the number of art objects offered for sale, the share of Modern and Contemporary items, and the top income. The variables are observed annually over the period 1976-2006. Each independent variable is expressed in log-difference. The price is the hedonic price index. Volume is the number of transactions observed each year based on data from Renneboog and Spaenjers (2013). Short-term transactions is the number of purchases that were resold within the next year, and come from the repeat-sale data set. Sales rate is the average percentage of items sold at auctions for each year. The number of art objects offered for sale is the proxy for the number of items offered at auctions, obtained by dividing the number of transactions by the sales rate. Modern and Contemporary indicates the share of Abstract Expressionism, Pop Art, and other Modern and Contemporary Art within the global art market (the aggregate of the thirteen art movements). Top income is the share of total income received by the top 0.1 percent of all income earners in the US, constructed by Piketty and Saez (2006). Coefficients significant at the 10% level are in bold.

Table IV: Price-Volume Correlation

	All Movements		High Volatility		Low Volatility	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Contemporaneous Correlation						
	Δ Price	Δ Price	Δ Price	Δ Price	Δ Price	Δ Price
Δ Volume	0.590***	0.544***	0.716***	0.704***	0.271**	0.247**
	(3.32)	(3.55)	(3.36)	(3.63)	(2.16)	(2.14)
Δ Stock		-0.133		0.036		-0.148
		(-0.67)		(0.15)		(-0.79)
Δ_{-1} Stock		0.284		0.075		0.241
		(1.41)		(0.32)		(1.29)
R^2	0.251	0.279	0.395	0.396	0.054	0.099
N	377	377	87	87	87	87
Panel B: Price Changes						
	(1)	(2)	(3)	(4)	(5)	(6)
	Δ Price	Δ Price	Δ Price	Δ Price	Δ Price	Δ Price
Δ_{-1} Price		-0.036		0.156		-0.056
		(-0.27)		(1.29)		(-0.42)
Δ_{-1} Volume	0.258**	0.235**	0.315**	0.216*	0.338*	0.312*
	(2.01)	(2.58)	(2.06)	(1.96)	(1.96)	(1.77)
Δ_{-1} Stock		0.318		0.321		0.210
		(1.20)		(0.86)		(0.59)
R^2	0.052	0.073	0.096	0.133	0.080	0.092
N	377	377	87	87	87	87
Panel C: Volume Changes						
	(1)	(2)	(3)	(4)	(5)	(6)
	Δ Volume	Δ Volume	Δ Volume	Δ Volume	Δ Volume	Δ Volume
Δ_{-1} Price	0.078	0.112	0.146	0.214*	-0.010	0.002
	(0.69)	(1.14)	(1.01)	(1.68)	(-0.09)	(0.02)
Δ_{-1} Volume		-0.160*		-0.162*		-0.179*
		(-1.82)		(-1.83)		(-1.85)
Δ_{-1} Stock		0.213		0.370		0.234
		(0.82)		(1.27)		(0.94)
R^2	0.007	0.037	0.021	0.062	0.000	0.036
N	377	377	87	87	87	87

Panel A and B report the estimation results for price changes on (lagged) volume changes, with and without controlling for lagged price changes and equity returns. Panel C comprises the estimation results for volume changes on lagged price changes, while controlling for lagged volume changes and equity returns. The annual series in Panel A are constructed from the full sample, while those in Panel B and C are constructed from fourth-quarter observations (see Section IV). We report the results for the aggregated art prices and volume (comprising 13 art movements), and for the aggregated three most volatile movements (Pop, Abstract Expressionism, Minimalism and Contemporary art) and the aggregated three least volatile art schools of art (Romanticism, Baroque, Rococo). Standard errors are clustered at year and movement level.
(t-Statistics are in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% level, respectively.)

Table V: Volume and Abnormal Returns

Decile	Purchase	Sale	Purchase	Sale	Purchase	Sale
	Full sample		High Volatility		Low Volatility	
Low	−0.99*	−8.25***	0.34	−3.75**	−0.52	−1.49
2	3.95***	−5.57***	5.51***	−12.92***	3.18	−4.21
3	2.24***	−3.44***	−1.62	−6.19***	0.50	−3.60
4	1.12**	−1.08***	4.30***	−4.40***	0.33	0.52
5	0.69	0.30	1.24	0.80	0.36	0.30
6	−0.27	0.13	0.60	−1.19	−1.93	0.72
7	0.07	1.02**	−4.23***	8.62***	−0.66	1.32
8	−2.29***	1.76***	0.84	5.92***	−0.09	3.34
9	−1.03**	5.90***	−4.38***	2.51*	−1.34	3.22
High	−3.52***	9.32***	−2.68**	10.82***	0.06	−0.12

This table presents the annualized returns in excess of the average (abnormal return) by volume decile (decile 10 corresponds to the largest volume). Each repeat-sale transaction is identified by purchase and sale date. We construct "portfolios" of paintings based on volume at the time of purchase or sale. Volume is constructed according to Equation (3.1). We compute the average abnormal return of each "portfolio" and rank them from the lowest volume (first decile) to highest (tenth decile) for the full sample and specific subgroups. Abnormal returns are defined as annualized returns in excess of the sample average. The "Full sample" corresponds to Figure 3.6. "High Volatility" are sales of Pop, Abstract Expressionist, and Modern and Contemporary art ($N = 2822$), and "Low Volatility" are sales of Romantic, Baroque and Rococo art ($N = 1107$). (***, **, and * indicate significance at the 1%, 5%, and 10% level, respectively.)

Table VI: Volume and Overpricing

	Full sample		High Volatility		Low Volatility	
	(1)	(2)	(3)	(4)	(5)	(6)
α	-0.001	-0.026***	0.013	-0.163***	-0.135***	-0.086**
	(-0.11)	(-3.55)	(0.69)	(-7.03)	(-4.76)	(-2.44)
MARKET	0.192***	0.332***	0.550***	0.569***	0.400***	0.433***
	(10.97)	(14.40)	(9.45)	(7.62)	(5.21)	(4.23)
SMB		0.353***		0.416***		-0.153
		(10.61)		(3.85)		(-1.02)
HML		-0.268***		-0.083		-0.727***
		(-9.07)		(-0.90)		(-5.18)
LIQ		0.169***		0.444***		0.475***
		(6.76)		(5.80)		(3.89)
FAME	0.147***	0.176***	0.129***	0.261***	0.008	0.001
	(13.05)	(15.17)	(2.91)	(5.95)	(0.16)	(0.02)
VOLUME	0.777***	0.751***	0.240**	0.432***	0.213*	0.183
	(29.10)	(27.97)	(2.43)	(4.38)	(1.73)	(1.51)
R^2	0.055	0.064	0.033	0.083	0.038	0.058
N	24889	24889	3361	3361	1052	1052

This table presents the estimates of the following regression:

$$\begin{aligned} \mathrm{r}_i - \sum_{t=b_i+1}^{s_i} r_{ft} &= \alpha + \beta \sum_{t=b_i+1}^{s_i} \mathrm{MKT}_t + \theta \sum_{t=b_i+1}^{s_i} \mathrm{SMB}_t + \phi \sum_{t=b_i+1}^{s_i} \mathrm{HML}_t \\ &+ \lambda \sum_{t=b_i+1}^{s_i} \mathrm{LIQ}_t + \gamma \sum_{t=b_i+1}^{s_i} \mathrm{FAME}_{a,t} + \nu \sum_{t=b_i+1}^{s_i} \mathrm{VOLUME}_{m,t} + \epsilon_i \end{aligned}$$

where $\mathrm{r}_i = \sum_{t=b_i+1}^{s_i} r_{it}$ is the return on item i between b_i and s_i, computed as the difference between the log of sale price and the log of purchase price and where r_{ft} is the risk free rate. The variable MKT_t is the world equity excess returns between purchase and sale times, SMB_t and HML_t are the Fama and French (1996) factors and LIQ_t is the Pastor and Stambaugh (2003) liquidity factor. $\mathrm{FAME}_{a,t}$ is the log of the share of mentions in Google Books for artist a at time t. $\mathrm{VOLUME}_{m,t}$ is the volume measure defined in Equation (3.1) for movement m. The three-stage-generalized-least square RSR estimation of Case and Shiller (1987) is used to estimate the regression for the three samples.
(t-Statistics are in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% level, respectively.)

Table VII: The "Pricing" of Volume by Subperiod

	1977-1986	1987-1996	1997-2006
CAPM	0.491***	1.070***	0.393***
	(11.89)	(29.58)	(4.91)
Fama-French	0.556***	0.981***	0.311***
	(13.26)	(26.92)	(3.87)

This table presents the estimates of the following regression:

$$\begin{aligned} \mathrm{r}_i - \sum_{t=b_i+1}^{s_i} r_{ft} &= \alpha + \beta \sum_{t=b_i+1}^{s_i} \mathrm{MKT}_t + \theta \sum_{t=b_i+1}^{s_i} \mathrm{SMB}_t + \phi \sum_{t=b_i+1}^{s_i} \mathrm{HML}_t \\ &+ \lambda \sum_{t=b_i+1}^{s_i} \mathrm{LIQ}_t + \gamma \sum_{t=b_i+1}^{s_i} \mathrm{FAME}_{a,t} \\ &+ \nu_1 \sum_{t=b_i+1}^{1986} \mathrm{VOLUME}_{m,t} + \nu_2 \sum_{1987}^{1996} \mathrm{VOLUME}_{m,t} + \nu_3 \sum_{1997}^{s_i} \mathrm{VOLUME}_{m,t} + \epsilon_i \end{aligned}$$

where $\mathrm{r}_i = \sum_{t=b_i+1}^{s_i} r_{it}$ is the return on item i between b_i and s_i, computed as the difference between the log of sale price and the log of purchase price and where r_{ft} is the risk free rate. The variable MKT_t is the world equity excess returns between purchase and sale times, SMB_t and HML_t are the Fama and French (1996) factors and LIQ_t is the Pastor and Stambaugh (2003) liquidity factor. $\mathrm{FAME}_{a,t}$ is the log of the share of mentions in Google Books for artist a at time t. $\mathrm{VOLUME}_{m,t}$ is the volume measure defined in Equation 3.1 for movement m. The variables ν_i, $i = 1 \ldots 3$ measure the impact of volume on excess returns for each sub-period: 1977-1986, 1987-1996 and 1997-2006. The three-stage-generalized-least square RSR estimation of Case and Shiller (1987) is used to estimate the regression for the three samples.
(t-Statistics are in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% level, respectively.)

Table VIII: Asymmetric Effects Of Lagged Returns On Volume

	All Movements (1)	High Volatility (2)	Low volatility (3)
	Δ Volume	Δ Volume	Δ Volume
Gains_{-1}	0.029	0.072	-0.133
	(0.14)	(0.25)	(-0.55)
Losses_{-1}	-0.184**	-0.329***	-0.098
	(-2.37)	(-3.57)	(-0.81)
Δ_{-1} Volume	-0.169**	-0.162*	-0.197**
	(-1.98)	(-1.90)	(-2.32)
Δ_{-1} Stock	0.272	0.466	0.325
	(1.04)	(1.58)	(1.09)
R^2	0.043	0.077	0.053
N	377	87	87

This table reports the estimated coefficients of a regression of volume on lagged returns, which separates positive from negative values of lagged log differences of art prices (gains and losses). The series are constructed from fourth-quarter observations (see Section IV.B and Table IV). The "Gains" (respectively "Losses") series correspond to change in price when the latter is positive (respectively negative) and zero otherwise. We report the results for the aggregated art prices and volume (comprising 13 art movements), and for the aggregated three most volatile movements (Pop, Abstract Expressionism, Minimalism and Contemporary art) and the three least volatile art paradigms (Romanticism, Baroque, Rococo). Standard errors are clustered at year and movement level.
(***, **, and * indicate significance at the 1%, 5%, and 10% level, respectively.)

Chapter 4

Return Predictability: Learning from the Cross-Section[1]

Chapter Abstract

This paper develops an estimation framework in which the true parameters of international return processes share a common distribution. The model (i) makes efficient use of the cross-sectional correlation in the residuals and (ii) learns about the common means and variances of the parameters. Once cross-sectional information is accounted for, the international evidence of return predictability appears much less fragmented than previously reported. In particular, there is reasonable evidence that the dividend-price ratio predicts both future returns and future dividend growth. Estimation risk is also substantially mitigated, so that stocks are typically safer in the long run.

I. Introduction

For about thirty years, researchers have been documenting the ability of various variables to forecast stock returns. This is not surprising because the existence of return predictability is not only of interest to practitioners but also has important bearing on the theoretical modeling of asset prices. In spite of a considerable literature in the last

[1]I am thankful for comments from Christophe Boucher, Edouard Challe, Guillaume Chevillon, Stéphane Chrétien, Pierre Collin-Dufresne, François Dezorme, Joost Driessen, Rik Frehen, Stéphane Gallon, Didier Janci, Frank de Jong, Malik Kerkour, Tim Kroencke, Martial Lasfargues, Isabelle Laudier, Lorenzo Naranjo, Patrice Poncet, Olivier Scaillet, as well as seminar participants at EPFL@SFI, Cambridge University, Ecole Polytechnique, ESSEC, THEMA, Caisse des Dépôts et Consignations and participants in the following conferences: AFFI 2014 (Aix), EFMA 2014 (Rome), ADRES Doctoral Conference 2014 (Paris) and Paris X Doctoral Conference in International Macroeconomics and Financial Econometrics. The support of the Institut Louis Bachelier-Collège de France Research Initiative "Long-Term Asset Allocation", together with its partners CNP Assurances and Caisse des Dépôts et Consignations is gratefully acknowledged. I am grateful to Benjamin Poignard (Institut Louis Bachelier) for exceptional research assistance. Part of this paper circulated previously as "International Return Predictability and the Term Structure of Risk." All remaining errors are my own.

decades, there is still considerable uncertainty on the magnitude, and pervasiveness of return predictability. As stressed by Pastor and Stambaugh (2012), even after observing more than two centuries of US data, "investors do not know the values of the parameters of the return-generating process, especially the parameters related to the conditional expected return." International evidence of return predictability is even less conclusive. If anything, the prevailing view seems to be that "return predictability is neither a uniform, nor a universal feature across international capital markets" (Schrimpf, 2010). There is also conflicting evidence on the relative importance of cash flow and discount rate news in explaining asset prices fluctuations. For example, Ang and Bekaert (2007) claim that "the dividend yields predictive power to forecast future dividend growth is not robust across sample periods or countries."

An important source of information has however been neglected in the previous literature: cross-sectional evidence of return predictability. Many studies, obviously, document evidence of return predictability outside the United States,[2] but most of these studies consider each country in isolation. The approach developed in this paper — cross-sectional learning[3] — considers them jointly. Two mechanism can improve precision by exploiting cross-country information. Consider a world with two countries:

$$\begin{aligned} r_{t+1}^{US} &= \mu^{US} + \beta^{US} x_t^{US} + u_{t+1}^{US} \\ r_{t+1}^{UK} &= \mu^{UK} + \beta^{UK} x_t^{UK} + u_{t+1}^{UK} \end{aligned}$$

where, in each country, excess returns are predictable by a variable x_t, e.g. the domestic dividend-price ratio. The first mechanism is to treat international data as a set of seemingly unrelated equations (Zellner, 1962). International stock returns are correlated, an information that is ignored in single-country regressions. The OLS estimator of the parameter of interest, say $\hat{\beta}^{US}$, is not efficient, because it assumes that the residual covariance is diagonal, i.e. $cov(u_{t+1}^{US}, u_{t+1}^{UK}) = 0$. An improved mechanism should exploit the covariance of the residuals.

[2]Recent studies include Campbell (2003),Rapach et al. (2005), Ang and Bekaert (2007), Driesprong et al. (2008), Schrimpf (2010), Rapach et al. (2013). Hjalmarsson (2010) provides the most comprehensive empirical investigation of stock return predictability worldwide, using a panel data that include over 20,000 monthly observations from 40 international markets.

[3]Jones and Shanken (2005) coined the term for the purpose of evaluating mutual fund performance.

The second mechanism is to treat the parameters of each country's return process as random variables. For example, suppose that I estimate $\hat{\beta}^{US}$ to be close to zero, but that at the same time I find a strong relation between the dividend-price ratio and future returns in the United Kingdom. The previous literature considers these two elements of evidence as distinct facts.[4] In the framework developed in this paper, I will assume that coefficients are drawn from a common normal distribution. This evidence would conduct me to revise my prior about return predictability in both countries. The reason for doing so originates from de Finetti's (1964) exchangeability assumption, further developed by Lindley and Smith (1972) for linear regression models: when there is not enough information to allow for precise estimation of an individual effect, it is natural to assume that the difference with other individual effects is the work of chance.

To better understand the concept of exchangeability, consider the return process in a hypothetical country, Zembla.[5] We have histories of stock returns and dividend-price ratios for other countries, but not Zembla. In the absence of any relevant information about Zembla's return process, our best guess would be $\beta^Z \sim \mathcal{N}(\bar{\beta}, \sigma_{\bar{\beta}})$, where $\bar{\beta}$ and $\sigma_{\bar{\beta}}$ are the hyperparameters characterizing the "population" of international return generating processes. If we *do* have histories of stock returns and dividend-price ratios for Zembla, should we overlook the information about $\bar{\beta}$ and $\sigma_{\bar{\beta}}$? The single-country estimate of $\hat{\beta}^Z$ may be imprecise, perhaps less precise than the common mean $\bar{\beta}$. In that case, intuition suggests a better estimate of β^Z would put considerable weight to the common mean. If, in contrast, the single-country estimate were more precise, it would be given more weight.

This is, together with the previous mechanism, the main ingredient of the Bayesian approach that I pursue in this paper. This approach has several advantages over frequentist alternatives such as the maximum likelihood approach. First, although I assume a non-informative prior for the common hyperparameters, the Bayesian estimate will typically differ from the classical frequentist estimator. The effect of the exchangeable prior is that the individual coefficients will be shrunk to the common mean, when the former are not

[4] Ang and Bekaert (2007), Hjalmarsson (2010) and Rapach et al. (2013) consider pooled regressions, i.e. take the polar viewpoint that the slope parameters are identical across countries.

[5]
> Old Zembla's fields where my gray stubble grows,
> And slaves make hay between my mouth and nose.

(From "Pale Fire" by Vladimir Nabokov.)

precisely estimated. Second, the Bayesian approach allows to incorporate economically motivated constraints in the estimation process. In particular, if return predictability reflects time-varying risk premia, the equity premium forecasts should always be positive.[6] Finally, the Bayesian approach allows estimation risk to enter in the decision process, for example in portfolio choice.

I reexamine the international evidence of stock return predictability using a large data set of fifteen countries. I concentrate on the "traditional" macro-financial predictors of stock returns: the dividend-price ratio, the short-term interest rate and the term spread. I also examine the ability of the dividend-price ratio to forecast dividend growth. The reason for doing so is that the present value identity links stock prices to future returns and future dividend growth. If the dividend-price ratio varies over time, it should forecast expected return, future dividends, or both (Lettau and Van Nieuwerburgh, 2008; Cochrane, 2008; Koijen and Van Binsbergen, 2010).

My main empirical findings can be summarized as follows. Cross-sectional learning substantially alters coefficient estimates. The effect on parameter precision is unambiguously strong. This gains in precision manifests both in and out-of-sample, where forecasts based on the exchangeable prior typically outperform forecasts based on models that ignore the cross-section. Importantly, the international heterogeneity of individual estimates is substantially blurred, in contrast to previous literature. In particular, the dividend-price ratio significantly predicts both future returns and dividends. The evidence is remarkably homogeneous across countries, except in the UK and the US where there is stronger evidence of return predictability than of dividend growth. Out-of-sample, the term spread is a robust predictor of excess returns (except in the United States), but the performance significantly degrades when return forecast are constrained to be positive, suggesting that the term spread may be related to some form of mispricing.

I also evaluate the economic significance of cross-sectional learning by comparing the term structure of risk faced by an investor with these priors. A number of studies have recently revisited the popular belief that stocks are safer in the long run. Barberis (2000), Stambaugh (1999) and Campbell and Viceira (2002, 2005) find that the per-period vari-

[6]See, e.g. Merton (1980); Fama (1991); Kothari and Shanken (1997); Schwert (2003); Campbell and Thompson (2008); Driesprong et al. (2008); Pettenuzzo et al. (2013).

ance of stocks substantially decreases with horizon. In contrast, several recent studies that propose more realistic models accounting for model uncertainty and estimation risk show that the predictive volatility increases with horizon (Pastor and Stambaugh, 2012; Pettenuzzo and Timmermann, 2011; Johannes et al., 2014). I find that estimation risk is substantially mitigated once cross-sectional information is accounted for, illustrating the importance of estimation precision on portfolio choice. When economically motivated restrictions are *a priori* assumed, the term structure is remarkably flat, except again in the Unites States and the United Kingdom.

The rest of this paper is organized as follows. Section II introduces the model with learning from the cross-section and provides an overview of the estimation procedure. Section III presents the empirical results, and illustrates the consequence of learning on the term structure of volatility. Section IV concludes.

II. Methodology

A. A model of international return predictability

It is common in the return predictability literature (e.g. Kandel and Stambaugh, 1996; Stambaugh, 1999; Barberis, 2000; Wachter and Warusawitharana, 2009) to use vector auto regressions (VAR) to capture the relation between asset returns and predictor variables. I follow this literature and assume that the returns of stocks in excess of the risk-free rate $r_{i,t}$ is a linear function of lagged predictors such as the dividend-price-ratio. For the sake of simplicity I will concentrate on a single regressor, $x_{i,t}$ that follows an AR(1) process. The model takes the form:

$$r_{i,t+1} = \mu_i + \beta_i x_{i,t} + u_{i,t+1} \tag{4.1}$$

$$x_{i,t+1} = \alpha_i + \rho_i x_{i,t} + v_{i,t+1} \tag{4.2}$$

where the innovations are normal, i.i.d. across t and cross-sectionally correlated with covariance Σ.

The model characterizes asset price dynamics in $i = 1, 2, \ldots N$ countries over time $t = 1, 2, \ldots T$. Taking expectation of Equation (4.1), we see that $\mu_i + \beta_i x_{i,t}$ is the conditional equity premium for country i. If β_i differs from zero, then the equity premium varies over time. A typical assumption of the literature is to impose that $|\rho| < 1$. In particular, present value models that impose transversality require that the dividend yield must be stationary. This implies that the regressor is stationary, although the value of ρ will be typically close to one.

It is helpful to rewrite the system (4.1) - (4.2) in stacked form:

$$\begin{bmatrix} r_{1,t+1} \\ x_{1,t+1} \\ \vdots \\ r_{N,t+1} \\ x_{N,t+1} \end{bmatrix} = \begin{bmatrix} 1 & x_{1,t} & 0 & & \cdots & & & 0 \\ 0 & & 1 & x_{1,t} & & & & \\ \vdots & & & & \ddots & & & \\ & & & & 1 & x_{N,t} & & \\ 0 & & & & & & 1 & x_{N,t} \end{bmatrix} \begin{bmatrix} (\mu_1, \beta_1)' \\ (\alpha_1, \rho_1)' \\ \vdots \\ (\mu_N, \beta_N)' \\ (\alpha_N, \rho_N)' \end{bmatrix} + \begin{bmatrix} u_{1,t+1} \\ v_{1,t+1} \\ \vdots \\ u_{N,t+1} \\ v_{N,t+1} \end{bmatrix}$$

or

$$y_{t+1} = X_t\theta + \epsilon_{t+1}, \ \ \epsilon_t \sim \mathcal{N}(0, \Sigma) \tag{4.3}$$

Most of the existing literature considers predictive regressions in isolation (or focus on a single country). Alternatively, several recent papers have considered pooled regression by assuming that the slope coefficient β_i are equal across countries (Ang and Bekaert, 2007; Hjalmarsson, 2010; Rapach et al., 2013). Intuitively, the first approach ignores the meaningful information contained in the cross-section, while the second makes the strong assumption that the data-generating processes are similar across countries. This paper takes another direction and considers Equation (4.3) as a system of seemingly unrelated regressions (SUR). If the error terms are correlated, it is well known that one can obtain more efficient estimates of the model by considering them jointly (Zellner, 1962).

Some care must be taken regarding the structure of the covariance matrix Σ. The vector of innovations ϵ_{t+1} is of dimension $2N$ and therefore requires to handle $2N(2N + 1)/2$ covariance parameters. Most of these parameters are likely to be redundant and imprecisely estimated. To circumvent this problem, I assume the following factor structure for the innovations:

$$u_{i,t+1} = \delta_i^u \bar{u}_{t+1} + \tilde{u}_{i,t+1} \quad (4.4)$$

$$v_{i,t+1} = \delta_i^v \bar{v}_{t+1} + \tilde{v}_{i,t+1} \quad (4.5)$$

The factors $\bar{u}_t$ and $\bar{v}_t$, are normal random variables with zero means, variances σ^u and σ^v and correlation coefficient γ. The idiosyncratic part of the innovations, $\tilde{u}_{i,t}$ and $\tilde{v}_{i,t}$, are also normal random variables with zero means and variances σ_i^u and σ_i^v. The idiosyncratic part are independent from the factors and across countries but can be correlated within a given country, i.e. $corr(\tilde{u}_{i,t}, \tilde{v}_{i,t}) = \gamma_i$ if $i = j$ and zero otherwise. This specification allows for country-specific variances and let the correlation between the unexpected returns $u_{i,t}$ and the innovation in the predictor $v_{i,t}$ vary across countries.

B. *Prior beliefs*

Equation (4.3) is agnostic about the joint nature of the individual countries coefficients. Economic theory offers little guidance as to why predictability should be high in some countries and low in others. As argued by Hjalmarsson (2010), countries that share many common characteristics should be more likely to exhibit similar predictability patterns than those that do not. Menzly et al. (2004) propose an external habit persistence model close to Campbell and Cochrane (1999), that generates cross-sectional heterogeneity in return predictability. This heterogeneity follows from the difference in asset cash flows' exposure to aggregate consumption. Although their model studies predictability across industries, the argument can be extended to international markets, viewing each country as an individual asset (Hjalmarsson, 2010). However Menzly et al. suggest that the heterogeneity is likely to be small.

Other plausible sources of international heterogeneity are heterogeneous preferences (e.g. risk aversion), institutional differences, and measurement errors. In particular, for the dividend-price ratio, Engsted and Pedersen (2010) and Chen et al. (2012) sketch two possible reasons for cross-country heterogeneity. First, the fraction of dividend-paying firms has substantially decreased in the postwar period, but at different paces. The US and the UK seem to be the exception rather than the rule, where the drop in dividend-paying firms has been more dramatic than in continental Europe (Fama and French, 2001; Von Eije and Megginson, 2008). Second, US firms have increasingly engaged in dividend smoothing in the postwar period, blurring the evidence of dividend growth predictability and artificially increasing the evidence of return predictability. Renneboog and Trojanowski (2007) and Andres et al. (2009) document that UK companies smooth dividends in much the same way as US companies, while German firms tend to have more flexible dividend policies.

The point of view of this paper is that the country data generating processes share some similarity. Formally, each country's vector of coefficient $\theta_i = (\mu_i, \beta_i, \alpha_i, \rho_i)'$ is assumed to be normally distributed around a common mean:

$$\theta_i = \bar{\theta} + \eta_i \tag{4.6}$$

where η is normal with mean zero and diagonal covariance matrix T. Such a prior is often denoted as *exchangeable* (de Finetti, 1964) and, provided individual coefficients are not too dispersed, causes the Bayesian estimates to be shrunk towards the common mean. As argued earlier, I further restrict the coefficient ρ to be between -1 and 1, which rules out nonstationary behavior of the predictor.

It would be of interest to refine this assumption further by enriching the right-hand-side by country-specific characteristics, or to identify clusters of countries sharing similar characteristics. Such analysis would typically require a larger number of countries. For example, Canova (2004) applies a similar framework to identify convergence clubs in income per capita growth, but uses a dataset of 144 European regions. Alternatively, one could increase N by studying return predictability at the industry level. I leave this possibility to further research and concentrate on international evidence, in order to facilitate comparison with the earlier literature.

I choose a prior that is uninformative in the sense of Jeffreys (1961) for the remaining parameters, namely Σ, $\bar{\theta}$ and T. This approach nests as special cases the traditional approach that treats each data-generating process independently, and the pooled approach where cross-country heterogeneity is assumed to be negligible. The former arises when one assumes that individual parameters have little in common and ignores the cross-sectional correlation of the innovations (i.e. when $\mathrm{T}^{-1} = 0$ and when the off-diagonal terms of Σ are dogmatically set to zero for $i \neq j$). The latter corresponds to the assumption that $\theta_i = \bar{\theta}$.

Comparison with related studies

Using international information is not the only way to obtain more precise estimates of predictive regression. A complementary way is to use economically motivated constraints or informative priors on the predictive regression coefficients. Several approaches have been advocated by the literature. Kandel and Stambaugh (1996) let the investor come with some informative knowledge about the true data-generating-process. Pastor and Stambaugh (2009) and Pastor and Stambaugh (2012) impose that the sign of the correlation between shocks to unexpected and expected returns is negative. Wachter and Warusawitharana (2009) develop a class of informative priors that assign a low probability to high R^2 in the predictive regression. Koijen and Van Binsbergen (2010) introduce nonlinear restrictions in a log-linearized present-value model. Campbell and Thompson (2008) and Pettenuzzo et al. (2013) constrain the equity premium to be positive and show that that this restriction significantly improves out-of-sample return forecasts.

In a recent paper, Pettenuzzo and Timmermann (2011) let the model parameters follow a meta distribution allowing for structural breaks in the data-generating process. The spirit of Equation (4.6), is the same, but instead of assuming random time-variations in the data-generating process, I concentrate on cross-country variations. This paper is also related to Jones and Shanken (2005), in the portfolio choice literature, who develop a similar cross-sectional scheme where investors learn about mutual fund skill. Indeed in the present paper, β plays a roughly similar role to the intercept in the mutual fund studies. An important difference is that in the present setting, the cross-sectional correlation in the residuals is significant, giving rise to a seemingly unrelated structure to the system of interest.

C. Bayesian estimation

The Bayesian framework that I introduce below makes use of the two intuitions that were just sketched. First, it uses the cross-sectional country correlation to increase posterior precision. Second, the exchangeability assumption partially shrinks the coefficients to the common mean. The model takes the form of a three-stages hierarchy. The first stage corresponds to the likelihood of the data conditioned on country-level parameters. The second stage corresponds to the distribution of these parameters. The third stage corresponds to the distribution of the common means of the parameters. This setup allows to specify diffuse priors in the third stage of the hierarchy, and thus let the data "speak" about country heterogeneity, given the assumed likelihood. In this subsection, I provide a brief overview of the estimation methodology. Additional details are given in Appendix .B.

Let $D \equiv \{r_1, \ldots, r_T, x_0, x_1, \ldots, x_T\}$ represent the data available to the investor at time T. A posterior density of the parameter $\Psi = (\theta, \Sigma, \bar{\theta}, \mathrm{T})$ is computed as

$$p(\Psi|D) \propto L(D|\theta, \Sigma)p(\Psi) \tag{4.7}$$

where $p(\Psi)$ denotes the prior density of the parameters and $L(D|\theta, \Sigma)$ is the likelihood function for the seemingly unrelated regression model. From results in, e.g., Zellner (1971), the likelihood is given by[7]

$$L(D|\theta, \Sigma) = |\Sigma|^{-T/2} \exp\left[-\frac{1}{2}\sum_{t=1}^{T}(y_t - X_t\theta)'\Sigma^{-1}(y_t - X_t\theta)\right]. \tag{4.8}$$

The posterior distribution of the model's coefficients can then be obtained by integrating out the hyperparameters from the joint posterior density (4.7), which can be done with a Gibbs sampler (see, e.g. Hsiao et al. (1998), Chib and Greenberg (1995a)). Gibbs sampling is an iterative Markov Chain Monte Carlo (MCMC) procedure for obtaining a sequence of observations which are approximated from a specified multivariate probability distribution. Starting from some arbitrary initial values of the parameters, it samples successively from the posterior distribution of each parameter, conditional on the values of the other parameters sampled in the latest iteration. For this posterior density, I use a four-block Gibbs sampler as in Chib and Greenberg (1995a) (see Appendix .B). The first two blocks correspond to the individual parameters θ_i and of the covariance matrix

[7] This likelihood conditions on the first observation x_0. See Stambaugh (1999) for a discussion of alternative priors and the use of unconditional likelihood.

Σ. The last two blocks correspond to the meta distribution of the individual parameters, $\bar{\theta}$ and T. In particular, the conditional distribution for the country-specific parameters θ is given by

$$\theta|\Sigma, \bar{\theta}, \Delta \sim \mathcal{N}(m_\theta, V_\theta) \times \rho \in (-1, 1) \tag{4.9}$$

where $\Delta^{-1} = I_N \otimes \mathrm{T}^{-1}$ and

$$m_\theta = V_\theta \left(\Delta^{-1} A_0 \bar{\theta} + \sum_{t=1}^{T} X_t' \Sigma^{-1} y_t \right) \tag{4.10}$$

$$V_\theta = \left(\Delta^{-1} + \sum_{t=1}^{T} X_t' \Sigma^{-1} X_t \right)^{-1}. \tag{4.11}$$

Equations (4.10) and (4.11) give us insight on how cross-country information is accounted for in the posterior distribution and, hence, it is worth discussing them in detail. Note, first, that the Bayes estimator of θ differs from the classical sampling estimator. In the classical setting it makes no sense to estimate the individual parameters, because they are treated as random variables.[8] Conditional on $|\rho| < 1$, the latter is the SUR estimator, given by

$$\left(\sum_{t=1}^{T} X_t' \hat{\Sigma}^{-1} X_t \right)^{-1} \left(\sum_{t=1}^{T} X_t' \hat{\Sigma}^{-1} y_t \right),$$

with $\hat{\Sigma}$ being a consistent estimate of Σ. The conditional Bayes estimator, m_θ, is a weighted average of the SUR estimate and the common mean $\bar{\theta}$. The weights are respectively proportional to Δ^{-1} and $\sum_{t=1}^{T} X_t' \Sigma^{-1}$. In other words, the Bayes estimator gives weights to both time-series and cross-sectional evidence, and use weights that correspond to the precision of the time-series and cross-sectional information. As a result, the Bayes estimator shrinks the estimates of country coefficients to the common mean $\bar{\beta}$. As the time dimension increases, more information about individual coefficients becomes available, the information contained in the cross-section hence becomes relatively marginal and gradually converges to the SUR estimate (see Hsiao and Pesaran, 2004).

This weighted average form is reminiscent of the common Bayes estimator with an informative prior, which is a weighted average of the OLS estimate and the prior mean.

[8] In contrast, the frequentist estimator of the common mean $\bar{\theta}$ is identical to the Bayesian estimator, conditional on Δ and Σ, see Hsiao and Pesaran (2004).

Conditionally on $\bar{\theta}$ and Δ^{-1}, the prior for θ is indeed

$$\theta \sim \mathcal{N}(\bar{\theta}, \Delta), \ \ \rho_i \in (-1, 1), i = 1, \ldots, N.$$

Conceptually, $\bar{\theta}$ can be seen as an informative prior with precision Δ^{-1}, but unlike a true prior, this information proceeds from the data. As noted in Jones and Shanken (2005), ignoring this information is tantamount to specifying a joint prior distribution in which the beliefs about θ_i are independent across countries.

D. Equity premium constraint

Most models that successfully account for return predictability propose mechanisms where risk or risk premia vary over time. Mechanisms that are able to generate predictable returns include time-varying relative risk aversion (e.g. Campbell and Cochrane (1999)), time-varying aggregate consumption risk (Bansal and Yaron, 2004; Bansal et al., 2012) and time-varying consumption disasters (Gabaix, 2008). Quoting Pettenuzzo et al. (2013), "it is difficult to imagine an equilibrium setting where risk-averse investors would hold stocks if their expected compensations were negative, and so this seems like a mild restriction." Pettenuzzo et al. therefore suggest to restrict the predictive regression by constraining the parameters μ_i and β_i so that the forecasts $\mu_i + \beta_i x_{i,t}$ are positive at all points in time.

Although a negative equity premium could theoretically arise if stocks hedge against other risk factors (Boudoukh et al., 1997; Pettenuzzo et al., 2013), negative return forecasts are generally interpreted as evidence of mispricing. For example, Fama (1991) notes that "there is no evidence that low D/P signals bursting bubbles, that is, negative expected stock returns." Kothari and Shanken (1997) make a similar point and show that an investor who assigns a 50% probability that expected returns are never negative comes away with a posterior probability of only 8% for the period 1926-1991 (although they conclude that they are always positive of the postwar period). Driesprong et al. (2008) finds that oil prices frequently forecasts negative stock returns, and argues against an interpretation in terms of risk premia.

It is therefore of theoretical interest to understand the impact of this restriction to the magnitude of the slope coefficients β and, in our case, on their cross-country dispersion. While Pettenuzzo et al. (2013) show that this restriction consistently improves the accuracy of equity returns out-of-sample, they tend to find weaker evidence of in-sample predictability. In particular, a Bayesian investor who believes that the equity premium cannot be negative would assign a much higher posterior probability to cases where the log-dividend-price ratio does *not* predict the equity premium. The reason for this apparent paradox is that the equity premium restriction puts an upper bound on the magnitude of return predictability. For example, when the dividend-price ratio is very low, a positive β_i will predict low returns. But β_i cannot be too large, otherwise the model could forecast negative returns.

Following Pettenuzzo et al. (2013), this restriction (thereafter equity premium constraint or EP) can be incorporated in this setting by modifying the priors on the intercept and slope coefficients to belong to the set $E_\theta = \{\mu_i + \beta_i x_{i,t} \geq 0, i = 1, \ldots, N, t = 1, \ldots, T\}$. The constraint is imposed in the first block of the Gibbs sampler. In order to draw from the restricted posterior distribution, I use the Metropolis-Hastings algorithm (see Chib and Greenberg, 1995b; Griffiths, 2003; Johannes and Polson, 2003). The reader is referred to Appendix .B for details.

III. Empirical Results

A. Data

My data set consists of quarterly data for fifteen OECD countries: Australia (AUS), Belgium (BEL), Canada (CAN), Denmark (DNK), France (FRA), Germany (DEU), Italy (ITA), Japan (JPN), the Netherlands (NLD), Norway (NOR), Spain (ESP), Sweden (SWE), Switzerland (CHE), the United Kingdom (GBR) and the United States (USA). Due to data availability, sample periods differ between countries, the longest being the United States (1953 to 2013) and the shortest being the Netherlands (1986 to 2013). More countries could be included, but the benefit of cross-sectional learning may be outweighed by country-specific estimation risk by using data available on shorter time periods. Appendix .A provides further detailed information on data sources and construction.

The dependent variables are log excess stocks returns, although I also consider the predictability of dividend growth in Section III.C. Following Chen (2009), I construct series of log dividend growth and log dividend price-ratio under the assumption that dividends are not reinvested, and that they are reinvested at the market rate. The dividend-price ratio (Rozeff, 1984; Fama and French, 1988; Campbell and Shiller, 1988) is one of the most popular predictor of stock returns. It is related to future stock returns via the present-value identity. I also consider two traditional predictors of excess returns documented since the 1970s:[9] the short rate (see, e.g. Fama and Schwert, 1977) and the term spread (see e.g. Keim and Stambaugh, 1986; Campbell and Shiller, 1991).

Table I gives the first and second moments of the data for each country. Except for the log dividend-price ratio, the sample statistics are in annualized, percentage units. Although the time periods differ from one country to another, it is worth noting that the equity premium roughly lies between 4% and 10% for most countries during these sample periods, as previously reported in, e.g. Ang and Bekaert (2007) and Dimson et al. (2008). The US is the most profitable and least volatile market of the sample, with an equity premium of 10.8% and an annual volatility of 15.3%. The standard deviations of the four series are quite similar across countries and roughly follow the limited discrepancies of sample averages.

[**Insert Table I about here**]

B. Regression results

Bayesian estimates of the β_i for each of the fifteen countries are represented in Figures 4.1 to 4.3. Each figure shows boxplots of the posterior distribution under the single-country benchmark, the exchangeable prior and under the exchangeable prior where equity premium forecasts are constrained to be positive. The center line of each boxplot indicates the median of the distribution, the box and the vertical lines respectively include 75% and 99% of the observations. Tables II and III supplement the boxplots by providing point estimates and standard deviations. Table II, in particular studies regression for the excess returns and dividend growth, and is discussed in detail in Subsection III.C.

[9] Henkel et al. (2011) provide a chronology for macroeconomic variables that have been documented to predict excess returns. The literature has considered many other predictors of stock returns, including the default premium (Keim and Stambaugh, 1986;Fama and French (1989)) and the consumption-wealth ratio (Lettau and Ludvigson, 2001). Unfortunately, these predictors are only available for a limited number of countries, or for short time periods, and therefore cannot considered.

[Insert Figures 4.1 to 4.3 and Tables II and III about here]

The top panel of each figure represents the estimates under the traditional framework that ignores cross-country information. For the dividend-price ratio and the term spread, the median coefficients are mostly positive (except Germany and Italy). Symmetrically, the coefficients associated to the short-term interest rate are all negative (except for Italy, Japan and Sweden). Finally, the coefficients for the term spread are mostly positive (except for Australia and Italy). Overall, the posterior distributions differ greatly across countries, as previously documented in the literature. Some predictors are clearly significant in some countries, but there are also a few "outliers" where the coefficient is the wrong sign or unrealistically large.[10] The dividend-price ratio appears as a quite significant predictor for Australia and the United Kingdom, as reported in, e.g., Hjalmarsson (2010). There is also some evidence of predictability for the short-term interest rate (see Ang and Bekaert, 2007; Rapach et al., 2013) and the term spread (Hjalmarsson, 2010), but the coefficients are estimated with low precision.

Is there enough evidence to support this large heterogeneity? Panels (b) of Figures 4.1 to 4.3 present the posterior distribution once cross-country information is accounted for. For comparison purposes, the classical SUR estimates are also reported (denoted by diamonds). The cross-country heterogeneity documented on Panels (a) is clearly blurred, first thanks to the cross-sectional correlation of the innovations (as can be seen from the SUR estimates), second thanks to the exchangeable prior, which further shrinks the coefficients to the cross country mean. This suggests that there are, by large, enough similarities to make a meaningful use of cross-country information. The gain in precision is visually striking. For example, we learn from Table II that the standard deviation of the posterior distribution for β_i is typically half smaller. Using information from the cross-section hence yields estimates that are typically twice more precise. Countries with limited data (e.g. the Netherlands) and countries where the benchmark estimates are statistically fragile (e.g. Japan, Norway) gain more from cross-sectional learning. The impact is also dramatic on the coefficient signs and magnitude. Cross-country learning weights down coefficients that are anomaly large and push up those that are too low or

[10] Ross (2005) and Zhou (2010) derive theoretical bounds on the economic magnitude of return predictability (see also Rapach and Zhou (2012)).

of the wrong sign with respect to other countries. There is reasonable evidence that the short-term interest rate negatively predict future returns, as for all countries at least 75% of the posterior distribution is less than zeros. This is also true for the log-dividend price ratio (except Italy), but not for the term spread.

Bayesian estimates of the β_i, when the prior is modified to include equity premium constraints, can be seen on Panels (c) of Figures 4.1 to 4.3. The gain in precision and the shrinkage of individual coefficients toward the common mean is further amplified (note the change in scale for the vertical axis). As noted by Pettenuzzo et al. (2013), the constraint has to hold at each point in time (and, in my case, for each country simultaneously), therefore the number of constraints grows in proportion with the length of the sample size. This large set of potentially binding constraints pins down the coefficients of the predictive regression, and thus increases precision.[11] As expected, the boxplots indicate that their economic significance is largely reduced. In the United States, a one standard deviation increase in the dividend-price ratio raises next-quarter returns by 0.55% (0.63%) in the benchmark (cross-sectional learning) case. Once equity premium restrictions are *a priori* assumed, the economic effect falls down to 0.32%. The effect is particularly strong for the term spread, where the typical impact is close to zero (with the exception of the UK, the US, and to a lesser extent, Denmark). This suggests that equity premium constraints may not be appropriate for the term spread, and casts some doubt on its relation with risk premia. Put differently, these results might indicate that the term spread is a better predictor of *returns* than of risk premia, pointing toward mispricing. I provide further evidence for this thesis in Section III.D.

As can be seen on Figures 4.6 to 4.8 of the Internet Appendix, the log-dividend price ratio is more persistent when the equity premium constraint is entertained, while the persistence coefficients of the remaining predictors are quite unchanged. Stambaugh (1999) shows that Bayesian inference can provide surprisingly strong evidence for predictability with respect to frequentist estimators. When the correlation between innovations in returns and innovations in the predictive variable is negative, underestimating the persistence ρ_i leads to overestimating the evidence for predictability β_i.[12] Typically, the Stam-

[11]A important difference is that economic constraints increase precision by shrinking the set of admissible coefficients, while the exchangeable prior use cross-country information to increase precision.

[12]Stambaugh (1999) shows that the relation between β and ρ is nearly linear under various specifications

baugh correlation is strongly negative for the dividend price ratio, close to zero for the short-term interest rate and slightly positive for the term spread.[13] For the dividend-price ratio, the EP constraint rejects draws of β_i that are too large, which also corresponds to draws of ρ_i that are too small. For the short-term interest rate and the term spread where the Stambaugh correlation is much smaller, the impact on the predictor's persistence is negligible.

C. *Does the dividend-price ratio predict dividend growth?*

There has been considerable debate in the recent literature regarding whether stock prices move because of changes in expected returns or because of changes in expected dividend growth. The present-value relationship implies that stock prices cannot move unexpectedly unless one or the other changes over time (otherwise the dividend-price ratio should be constant over time; see, e.g. Cochrane, 2008). Because neither expected returns nor expected dividend growth can be observed, the debate centers on whether the dividend-price ratio predicts (excess) returns, dividend growth, or both.

The prevailing view is that stock returns are predictable, while dividend growth is not predictable, so that almost all variation in the price-dividend ratio comes from variations in expected returns (see e.g. John Cochrane's presidential address (Cochrane, 2011)). Several recent studies, however, have criticized this view on several ground.[14] For example, Chen (2009) and Koijen and Van Binsbergen (2010) show that it is sensitive to the assumption about the reinvestment rate of dividends received within the year. Although the present-value relation implies that dividends must be reinvested at the stock market rate (Cochrane, 2008), doing so is problematic because it imparts some of the properties of returns to dividends, strengthening the evidence of return predictability. The predictive

and proves that the relation is given by

$$E(\beta_i|D) - \hat{\beta}_i \approx E\left(\frac{\sigma_{i,uv}}{\sigma_v^2}\middle|D\right)[E(\rho_i|D) - \hat{\rho}_i] \tag{4.12}$$

where $\hat{\beta}_i$ and $\hat{\rho}$ are the OLS estimates of β_i and ρ_i, and $\sigma_{i,uv}$ and σ_v^2 are elements of the covariance matrix Σ. $\sigma_{i,uv}$ is often referred to as the Stambaugh correlation.

[13]The posterior mean for the Stambaugh correlation is respectively -0.76, -0.13 and 0.13 for the United States data.

[14]Koijen and Van Nieuwerburgh (2011) surveys the recent literature return and cash flow growth predictability.

relation also appears to be unstable, both in the time series and the cross-section dimension. Chen (2009) argues that the return predictability in the United States is essentially a postwar phenomenon. Maio and Santa-Clara (2012) show that dividend growth is predictable for portfolios of small and value stocks. Engsted and Pedersen (2010) provide similar evidence for Sweden and Denmark.

Table II presents Bayesian estimates for the regression $z_{i,t+1} = \mu_i + \beta_i x_{i,t} + u_{i,t+1}$ where $z_{i,t}$ is excess return or dividend growth rate in quarter t, with and without reinvestment, and where $x_{i,t}$ is the log-dividend price ratio. The first two columns correspond the estimates for returns without reinvestment, also reported in Panels (b) and (c) of Figure 4.1. As expected, we see that the evidence of return predictability is considerably stronger, economically and statistically, when dividends are reinvested at the market rate.[15]

Can the dividend-price ratio predict dividend growth? Again, Table II illustrates the importance of reinvestment assumption. The last column corresponds to the prevailing view: when dividends are reinvested at the market rate, there is no evidence that dividend growth is predictable. The typical distribution of β_i is remarkably centered about zero for all countries, except for the UK where there is some evidence of predictability. Without reinvestment, however, the effect reverses: there is considerable evidence of dividend growth predictability for all countries. The evidence is economically the weakest in the UK and the US, as previously documented by Engsted and Pedersen (2010) (who use annual data on a longer time span).[16]

In summary, there is substantial evidence that both excess returns and dividend growth are predictable. In the US and the UK, stock returns are strongly predictable, as previously noted in the analysis of the equity premium constraint, and dividend growth is modestly predictable. The pattern is reversed in the remaining countries, where return predictability is more modest and dividend growth predictability is pervasive, unless dividends are reinvested at the market rate.

[15]This is also true when the cross-sectional information is neglected (not shown). In particular, the coefficients are of the "right" sign for all countries

[16]I also estimate the same regressions with annual (non-overlapping) data. The results, reported in Table V of the Internet Appendix, confirm that dividend growth is strongly predictable, even in the US and UK.

D. Out-of-sample performance

The previous sections show how cross-sectional information can be exploited to obtain more precise estimates of predictive regressions. The results indicate that, under the relatively mild prior that countries share some similarities, international predictability is much less fragmented that would indicate the estimation of individual predictive regression, and that estimation risk is strongly reduced. In this section, I evaluate the out-of-sample performance of this prior. Out-of-sample evidence is increasingly viewed an essential criterion to assess stock return predictability. In particular, Welch and Goyal (2008) show that a simple forecasting rule based on the historical average of past returns outperformed most of predictors that had been suggested by the literature, casting doubt on the evidence that risk premia are time-varying.

A number of papers has in turn proposed more sophisticated models that can improve the statistical and economic value of out-of-sample forecasts. A common argument is that estimation on short samples, combined with model uncertainty and parameter instability, render conventional predictive regression forecasts unreliable. Hence, imposing reasonable restrictions on the forecasts or coefficients estimates can reduce forecast errors.[17] As previously mentioned in Section II.D, Campbell and Thompson (2008) and Pettenuzzo et al. (2013) impose that the equity risk premium must be positive, and obtain out-of-sample forecasts that are better than the unconditional mean, which implies that investors could have profited by using market-timing strategies.[18]

From a forecasting point of view, the approach advocated in this paper can be seen as further imposing sensible restrictions to coefficient estimates. Pooling models have been successfully used for the purpose of GDP growth forecasting (see, e.g. Mittnik, 1990; Zellner and Hong, 1989; Hoogstrate et al., 2000). Hjalmarsson (2010) finds that equity premium forecasts based on the fixed effect estimator often outperform those based on

[17] In fact, even wrong restrictions can improve out-of-sample forecasts. This occurs when the bias resulting from imposing false restrictions is outweighed by the reduction of the estimator's variance due to the restriction (see e.g. Hoogstrate et al., 2000).

[18] Several other studies impose economically motivated constraints, as discussed in Section II.D. Many other approaches have been shown to successfully improve out-of-sample forecasts of equity returns, including forecast combinations and Bayesian model averaging (see, e.g., Cremers, 2002, Avramov, 2002, Rapach et al., 2009, Schrimpf, 2010, Dangl and Halling, 2012), factor models (e.g. Ludvigson and Ng, 2007, Kelly and Pruitt, 2012, Neely et al., 2014) and regime and time-varying coefficient models (e.g. Pesaran and Timmermann (2002), Paye and Timmermann, 2006, Rapach, 2006, Henkel et al., 2011, Pettenuzzo and Timmermann, 2011, Dangl and Halling, 2012, Johannes et al., 2014).

the time-series estimates. This paper goes one step further by allowing partial pooling and by incorporating risk premium restrictions in the Bayesian estimation.

I evaluate out-of-sample predictive ability by comparing forecasts based on predictive regressions and forecasts based on historical average of country excess returns. As shown by Welch and Goyal (2008), the historical average is a stringent benchmark: forecasts based on macroeconomic variables rarely achieve to outperform the historical average forecast out of sample. Following Hjalmarsson (2010), I exclude countries with less than 40 years of data. I use a forecasting period of 20 years, which leaves minimally 20 years of in-sample training period for each country. The forecasting exercise occurs recursively, on an expanding window. For each country, I compare historical average forecasts to forecasts generated from a competing predictive regression model. The historical average forecast corresponds to the constant expected excess return model, i.e. $\hat{r}^i_{t+1} = \hat{\mu}_{0i,t}$, while the predictive regressions forecasts are obtained as $\hat{r}^i_{t+1} = \hat{\mu}_{i,t} + \hat{\beta}'_{i,t} z^i_t$. Forming forecasts in this manner simulates the situation of an investor in real time.

To compare the historical average forecasts against the competing predictive regression forecasts, I use Campbell and Thompson (2008) out-of-sample R^2 out-of-sample statistic, R^2_{OS}. The R^2_{OS} measures the proportional reduction in mean-squared forecast error (MSFE) for the competing model relative to the historical average benchmark:

$$R^2_{OS} = 1 - \frac{\left(\sum_{t=1}^{T} r^i_t - \hat{r}^i_t\right)^2}{\left(\sum_{t=1}^{T} r^i_t - \bar{r}^i_t\right)^2} \tag{4.13}$$

where $\hat{r}^i_t$ and $\bar{r}^i_t$ are respectively forecasts obtained by predictive regression and forecasts based on the historical mean. A positive R^2_{OS} implies that the predictive regression has lower average mean-squared prediction error than the historical average return. I also compute the Clark and West (2007) *MSFE-adjusted* statistic to test the null that the historical average MSFE is less than or equal to the predictive regression MSFE, against the alternative hypothesis that the historical average MSFE is greater than the predictive regression MSFE. This test is one-sided and corresponds to $H_0 : R^2_{OS} \leq 0$ against $H_A : R^2_{OS} > 0$. The Clark and West (2007) test is itself a modification of the Diebold and Mariano (1995) and West (2006) statistic. The modification ensures that asymptotically the statistic approximately follows a standard normal distribution, when comparing forecasts from nested models. In order to compare least-square forecasts from the forecasts obtained with the Bayesian framework, i.e. to compare non-nested models, I use the Diebold and Mariano (1995) and West (2006) statistic (DMW). This test corresponds to $H_0 : MSFE_{Bayes} \leq MSFE_{LS}$ against $H_A : MSFE_{Bayes} > MSFE_{LS}$.[19]

Table IV reports the out-of-sample results for the baseline least-squares and Bayesian forecasts. I use the same set of predictors as in the previous section; each panel considers the performance of a given predictor using least-squares and Bayes coefficients. For each of these three models, I also consider forecast obtained after imposing the equity premium constraint introduced in Section II.D.[20] The first group of three columns shows out-of-sample results without imposing equity premium constraints on the coefficients and the second group impose equity premium constraints. Within each group, the table presents the R^2_{OS} for the two models, as well as the difference between the two R^2_{OS}. The table also indicates whether the reported results are statistically significant.

[Insert Table IV about here]

The first striking result of Table IV is the remarkable improvement in out-of-sample forecasts. Although the primary goal of the present paper is not to generate out-of-sample performances, the cross-sectional approach generates — on average — superior forecasts (although the absolute performance is overall mixed, as I discuss below). Learning typically reduces large forecasting mistakes. For example, for the US when forecasting with the dividend-price ratio, the R^2_{OS} raises from -4.24% to -1.71%. We see from Table IV that economically constrained models also tend to produce better return forecasts than unconstrained forecasts. However, the typical improvement is too small to consistently outperform the historical average. Interestingly, the benefit of learning remains once EP restrictions are imposed. Forecasts with both EP restriction and learning slightly outperform (again on average) the historical benchmark.

[19]It is tempting to say that this is equivalent to $H_0 : R^2_{OS,Bayes} < R^2_{OS,LS}$ against $H_A : R^2_{OS,Bayes} > R^2_{OS,LS}$. A subtle difference is that R^2_{OS} is computed from a ratio of MSFE, while the DMW is computed as a difference. Therefore, it is possible that a statistically superior model (according to the DMW statistic) obtains a lower R^2_{OS} than a statistically inferior model.

[20]For the baseline model I do not implement a Bayesian model as in Pettenuzzo et al. (2013), but rather obtain estimates using nonlinear least squares. While both methods yield qualitatively similar forecasts, the latter is computationally faster.

Turning to the performance of individual predictors, the results in panel A of Table IV highlight the poor out-of-sample predictive power of the dividend-price ratio. Only one of the R^2_{OS} (Japan) is positive in the baseline case and none is significant. This poor predictive power is in line with the previous literature and motivated more sophisticated forecasting models. The average R^2_{OS} is -2.88% for the baseline model and -1.27% for the Bayesian approach, and 7 countries out of 11 gain from the latter.

Panel B and C report the out-of-sample R^2 for the short-term interest rate and the term spread. The Bayesian approach consistently outperforms least-squares for both predictors, the "best" model being, by far, the cross-sectional model using the term spread as predictor. Under this latter specification, the out-of-sample forecasts consistently beat the historical average for nine countries (significantly in five), the average R^2_{OS} being 1.21%.

Overall, among the traditional variables used to predict returns, only the term spread performs consistently well out-of-sample, in line with previous evidence. In particular, Hjalmarsson (2010) notes that forecasts based on pooled estimates generally outperform forecasts based on single-country estimates.[21] Erik Hjalmarsson also finds evidence of out-of-sample predictability from the term spread, for 10 out of 14 countries.

It is important to emphasize, however, that this out-of-sample exercise is conducted on a relatively short timespan (1993-2013). The previous literature typically entertains longer forecasting period (but concentrates on the US) and finds that returns were remarkably hard to predict in the last two decades. For example, Pettenuzzo et al. (2013), find that constrained forecasts of US equity returns outperform the historical average during the 1947-2010 period, but underperform during the 1979-2010 period. In a similar vein, the time-varying parameter model of Dangl and Halling (2012) fails to consistently beat the historical average during the 1988-2002 period. This is also true for studies that consider international evidence (see e.g. Henkel et al., 2011, Schrimpf, 2010, Hjalmarsson, 2010).

Finally, it is interesting to note that EP restrictions degrade performance when forecasting with the term spread. This evidence corroborates the in-sample results previously discussed and gives credit to the thesis that the term spread may reflect mispricing and

[21] I also considered pooled forecasts based on fixed effect estimates, as in Hjalmarsson (2010). The results (not shown) are typically inferior for the log-dividend price ratio and the short-term interest rate, and marginally better for the term spread. The tables are available upon request.

not risk.

E. The term structure of risk

In this final section, I study the consequence of cross-sectional learning on the term structure of stock volatility. Whether stocks are safer in the long run has received considerable attention in the recent literature, in particular following the thought-provoking paper of Pastor and Stambaugh (2012), who claim that stocks may be riskier in the long run.[22] An important component of long-term investment is estimation risk, and it is therefore worthwhile to study if cross-sectional learning makes stocks safer or riskier in the long run. We are interested in the following quantity:

$$\sigma_i^2(k|D) = \frac{1}{k} Var(r^i_{t\to t+k}|D) \tag{4.14}$$

where $r^i_{t\to t+k} = \sum_{j=1}^{k} r_{t+j}$ is the cumulated k-quarter ahead excess return. The predictive system (4.1)-(4.2) constitutes a restricted VAR that is commonly studied in the literature. To facilitate comparison with previous works, it is useful to rewrite it as:

$$\begin{pmatrix} r_{i,t+1} \\ x_{i,t+1} \end{pmatrix} = \begin{pmatrix} \mu_i \\ \alpha_i \end{pmatrix} + \begin{pmatrix} 0 & \beta_i \\ 0 & \rho_i \end{pmatrix} \begin{pmatrix} r_{i,t} \\ x_{i,t} \end{pmatrix} + \begin{pmatrix} u_{i,t+1} \\ v_{i,t+1} \end{pmatrix} \tag{4.15}$$

or more compactly

$$\mathbf{z}_t^i = \Phi_0^i + \Phi_1^i \mathbf{z}_{t-1}^i + \mathbf{v}_t^i \tag{4.16}$$

where $\mathbf{v}_t^i \sim \mathcal{N}(0, \Sigma_i)$.[23] Campbell and Viceira (2002, 2005) show that the VAR framework provides a natural setup to compute multi-period moments. They find that the conditional variance of stock return grows more slowly with the investment horizon, consistently with the "conventional wisdom" that stocks are safer in the long run, popularized by Siegel (1998). To see that, note that under the assumption that Σ_i is constant over time, the

[22] A partial list includes Siegel (1998); Stambaugh (1999); Barberis (2000); Campbell and Viceira (2002, 2005); Bec and Gollier (2009); Jondeau and Rockinger (2010); Favero and Tamoni (2010); Diris (2011); Pettenuzzo and Timmermann (2011); Cales et al. (2013); Hoevenaars et al. (2014), and Johannes et al. (2014).

[23] Therefore Σ_i are block diagonal elements of the $2N \times 2N$ matrix Σ.

conditional k-quarter variance is given by:[24]

$$\sigma_i^2(k|D,\Psi_i) = \frac{1}{k}\mathbf{M}\,Var\left(\mathbf{z}_{t+1}^i + \mathbf{z}_{t+2}^i + ... + \mathbf{z}_{t+k}^i\right)\,\mathbf{M}' \tag{4.17}$$

where the vector $\mathbf{M} = (1,0)'$ extracts returns from the vector $\mathbf{z}_t^i$; Ψ_i is the set of parameters for country i and

$$\begin{aligned} Var\left(\mathbf{z}_{t+1}^i + \mathbf{z}_{t+2}^i + ... + \mathbf{z}_{t+k}^i\right) &= \Sigma_i + \left(I + \Phi_1^i\right)\Sigma_i\left(I + \Phi_1^i\right)' \\ &+ \left(I + \Phi_1^i + \Phi_1^i\Phi_1^i\right)\Sigma_i\left(I + \Phi_1^i + \Phi_1^i\Phi_1^i\right)' \\ &+ ... \\ &+ \left(I + \Phi_1^i + ... + (\Phi_1^i)^{k-1}\right)\Sigma_i\left(I + \Phi_1^i + ... + (\Phi_1^i)^{k-1}\right)' \end{aligned} \tag{4.18}$$

Equation (4.17) describes the term structure of risk faced by an investor who understands that a fraction of stock returns is predictable. Perhaps counterintuitively, although returns are predictable, their volatility does not necessarily decreases with horizon. Consider the two-period volatility:

$$\frac{1}{2}Var(r_{t\to t+2}) = \frac{1}{2}Var(r_{t+1}) + \frac{1}{2}Var(r_{t+2}) + cov(r_{t+1}, r_{t+2})$$

Absent return predictability, returns will not be autocorrelated and the two-period variance will equal the sum of single-period variance, which we assume equal. If returns are predictable, the covariance term may be positive or negative, so that volatility may increase or decrease with the investment horizon. More generally, return dynamics will be captured by the VAR. An investor who neglects predictability, and therefore consider Equation (4.16) assuming that the coefficients of Φ_1^i are zero, will face a flat term structure of volatility.

Importantly, Equation (4.17) provides the term structure of risk conditional to a given set of parameters. Since a typical investor would be uncertain about the true values of these parameters, a Bayesian investor would rather be interested in the predictive variance

[24] See the Internet Appendix.

(Pastor and Stambaugh, 2012):

$$\sigma_i^2(k|D) = E\left(\sigma_i^2(k|D,\Psi_i)\big|D\right) + Var\left(E(r^i_{t\to t+k}|\Psi_i, D)\big|D\right) \tag{4.19}$$

Put differently, Equation (4.19) accounts for parameter uncertainty, while (4.17) does not. The second term of (4.19) is the variance of the conditional mean and adds positively to the expected conditional variance. The predictive variance will thus generally be larger than the conditional variance studied by Campbell and Viceira (2002, 2005).

Figures 4.4 and 4.5 plot the predictive volatility for each of the fifteen countries, with investment horizon increasing from 1 quarter to 15 years, for the three specifications introduced earlier. I concentrate on a model where returns are predicted by the log dividend-price ratio. The curves with continuous line correspond to the traditional case that ignores cross-sectional learning. We observe that for most countries, the predictive variance is larger at higher horizons. In this simple setup with a single predictor and where the VAR is restricted, mean reversion of returns is a byproduct of the correlation between innovations in returns and innovations in the predictive variable (the Stambaugh correlation $\sigma_{i,uv}$), and the predictive slope β_i (Campbell and Viceira, 2005). As stressed earlier, the correlation is strongly negative and the predictive slope is typically positive, inducing mean reversion in returns. This pertains to the conditional variance (the first term in Equation (4.19)). Once parameter uncertainty is accounted for, however, mean reversion is more than compensated by uncertainty about expected returns (the second term in Equation (4.19)). The effect becomes stronger with investment horizon, increasing the slope of the term structure. This is typically the case for countries where data is not available on a long history, such as the Netherlands, or where the predictive slope is estimated with low precision, such as Canada. For those countries, the predictive volatility is J-shaped and the 15-year predictive volatility can be substantially larger than the one-quarter volatility (e.g. by about 5% for the Netherlands). These results are largely in line with previous literature (see Bec and Gollier (2009) for French data and Jondeau and Rockinger (2010) for a larger set of European countries). In contrast, the predictive volatility is clearly downward slopping for the United States, as is typically the case in the literature with diffuse priors (Stambaugh, 1999; Barberis, 2000).

[Insert Figures 4.4 and 4.5 about here]

The predictive volatilities with the exchangeable prior are represented by dotted lines. We see that the curves are now downward slopping for most countries, reflecting two effects. First, the slope coefficients are partially shrunk toward the common mean, increasing mean reversions in countries where it was initially weak or negative. This is the case, e.g., for France and Germany. Second, parameter uncertainty is clearly mitigated, reducing volatility at long horizons. The consequence are of first order for most countries, the effect being a flattening of the term structures of risk, on average modestly decreasing by 1%-3%. Two notable exceptions are the US and the UK, where stocks remain manifestly safer in the long run, resulting from the larger sample size and remarkable precision of the baseline model. Finally, the predictive volatilities with the exchangeable prior and economic constraints are represented by dashed lines. I showed in Section III.B that imposing economic constraints tends to reduce the economic significance of predictability. We see on Figures 4.4 and 4.5 that the consequence on the term structure of risk is non negligible. The term structure of risk is univocally flat for all countries, again with the notable exception of the US and the UK.

IV. Conclusion

The previous literature has produced a number of comprehensive studies of international return predictability. This paper argues that these previous studies, by treating each country separately, have neglected important information about predictability as a whole. Two ingredients have been previously overlooked: (i) the contemporaneous correlation between international returns, or more precisely, the seemingly unrelated structure of international data; (ii) the joint nature of international return data generating processes.

I develop an estimation framework that assumes that the true parameters of international return processes share a common normal distribution. By treating the evidence of return predictability jointly, the researcher makes efficient use of the cross-sectional correlation of the data and learns about the common means and variances of the parameters. He can then use both country-specific and international estimates and weight them according to their respective precision. The resulting model nests as special case the classical approach that considers individual countries separately.

Departing from the traditional approach yields rich consequence on the evidence of return predictability. Cross-country heterogeneity appears much smaller than previously reported and, as a result, the model makes an efficient use of the joint information. In particular, there is substantial evidence that the dividend-price ratio forecasts both future returns and dividend growth. Learning also substantially alleviates estimation risk, and hence makes stocks typically safer in the long run.

Appendix

Appendix A. Data description

The data used in this paper are quarterly time series and come from four different sources: the CRSP, the Morgan Stanley Capital International (MSCI) database, the OECD, the IMFs International Financial Statistics and the St. Louis Fed FRED database. The following 15 countries are analyzed: Australia (AUS), Belgium (BEL), Canada (CAN), Denmark (DNK), France (FRA), Germany (DEU), Italy (ITA), Japan (JPN), the Netherlands (NLD), Norway (NOR), Spain (ESP), Sweden (SWE), Switzerland (CHE), the United Kingdom (GBR) and the United States (USA).

Monthly short-term interest rates are downloaded from the OECD statistics or FRED (USA) the IMF International Financial Statistics depending on availability. End-of-quarter values from this monthly series are retained to get quarterly observations (Y_t). y_t denotes the 3-month log yield: $y_t = log(1+Y_t)/4$. I construct term spreads (spr) from the IMF (FRED for the United States) monthly series of 10-year Government bond yields. I retain end-of-quarter values to get quarterly series. The term spread spr is defined as the difference between the long and short interest rate yields.

My US stock data is the Standard & Poor's Composite Index, downloaded from the CRSP. International equity prices and total returns (in local currency) come from Morgan Stanley Capital International (MSCI) database. Following convention, a smoothed dividend series (a sum of dividends from month $t-11$ through month t) is used to compute the dividend-price ratio. I also construct a series of dividends reinvested at the stock market rate, following Chen (2009), and obtain the series of log dividend growth g and g^r (the subscript indicates that dividends are reinvested). The log dividend-price ratios, dp and dp^r obtains as the log dividend less the log price index.

Appendix B. Bayesian framework

Likelihood and prior beliefs

It is convenient to rewrite Equation (4.3) in stacked form:

$$y = X\theta + \epsilon,\ \epsilon \sim \mathcal{N}(0, \Sigma \otimes I_T) \tag{20}$$

where $\theta = (\theta_1, \theta_2, \ldots, \theta_N)$, $y = (y_1', y_2', \ldots, y_N')'$, $\epsilon = (\epsilon_1', \epsilon_2', \ldots, \epsilon_N')'$, y_i and ϵ_i are $2T \times 1$ vector of the left-hand side variables and innovation terms for country i. Finally $X = \text{diag}(X_1, X_2, \ldots, X_N)$, $X_i = (\iota_T, (x_{i,1}, x_{i,2}, \ldots, x_{i,T})')$ where ι_T denotes a vector of ones.

Equation (4.6), which corresponds to the second stage of the hierarchy, can be rewritten as

$$\theta = A\bar{\theta} + \eta,\ \eta \sim \mathcal{N}_{NG}(0, \Delta),\ \Delta^{-1} = I_N \otimes \mathrm{T}^{-1} \tag{21}$$

where $A = (I_G, I_G, \cdots I_G)'$ maps the $G \times 1$ vector of common coefficients $\bar{\theta}$ to the $NG \times 1$ vector θ. $G = 4$ is the number of coefficients per country and $\otimes$ denotes the Kronecker product. Equation (21) says that the individual coefficients $\theta_i = (\mu_i, \beta_i, \alpha_i, \rho_i)'$ follow a normal distribution with $G \times G$ covariance matrix T.

The likelihood function of Equation (20) is written as follows:

$$p(D|\theta, \Sigma) \propto |\Sigma|^{-T/2} \exp\left\{-\frac{1}{2}(y - X\theta)'\left(\Sigma^{-1} \otimes I_T\right)(y - X\theta)\right\} \tag{22}$$

Further, I assume the following priors for the model parameters:

$$\begin{aligned}
p(\theta|\bar{\theta}, \Delta) &\propto |\Delta|^{-N/2} \exp\left\{-\frac{1}{2}(\theta - A_0\bar{\theta})'\Delta^{-1}(\theta - A_0\bar{\theta})\right\},\ \rho \in (-1, 1) \\
p(\Sigma^{-1}) &\propto |\Sigma^{-1}|^{-(NG+1)/2} \\
p(\bar{\theta}, \mathrm{T}^{-1}) &\propto |\mathrm{T}^{-1}|^{-(G+1)/2},\ \Delta^{-1} = I_N \otimes \mathrm{T}^{-1}
\end{aligned}$$

The prior on θ corresponds to the assumption that individual coefficients are drawn from a common distribution. For the remaining parameters, the priors are uninformative in the sense of Jeffreys (1961). I assume that the parameter vectors are mutually independent.

The joint posterior for $\Psi = \{\theta, \bar{\theta}, \Sigma^{-1}, \Delta^{-1}\}$ is thus

$$p(\Psi|D) \propto p(D|\theta, \Sigma)p(\theta|\bar{\theta}, \Delta^{-1})p(\bar{\theta})p(\Sigma^{-1})p(\Delta^{-1}) \tag{23}$$

Sampling from the posterior

The Gibbs sampler consists in four blocks. I initialize the Gibbs sampler with SUR estimates. I start the sampler using 2,000 draws that I discard; I use the subsequent sample of 10,000 draws for the purpose of inference in the main section (1,000 for the out-of-sample forecasts). When the equity premium constraint is entertained, the draws are strongly autocorrelated and keeping every tenth draws appears to be optimal for storage considerations. The convergence is supported by visual inspection of the posterior draws and by the MCMC diagnostics of Raftery and Lewis (1995, 1992b,a), Geweke (1992) numerical standard errors and relative numerical efficiency estimates, and the Geweke chi-squared test comparing the means from the first and last part of the sample.[25] I derive below the conditional distribution for each of the four parameters.

1. Conditional posterior for θ

Viewing the joint posterior in Equation (23) as a function of only θ yields the following conditional posterior for θ (conditional on $\rho \in (-1, 1)$):

$$\begin{aligned} p(\theta|D, \bar{\theta}, \Delta, \Sigma) &\propto p(D|\theta, \Sigma)p(\theta|\bar{\theta}, \Delta^{-1}) \\ &\propto \exp\left\{-\frac{1}{2}(y - X\theta)'\left(\Sigma^{-1} \otimes I_T\right)(y - X\theta)\right\} \exp\left\{-\frac{1}{2}(\theta - A_0\bar{\theta})'\Delta^{-1}(\theta - A_0\bar{\theta})\right\} \end{aligned}$$

Note that

$$(y - X\theta)'\left(\Sigma^{-1} \otimes I_T\right)(y - X\theta) = (\theta - \hat{\theta})'X'\left(\Sigma^{-1} \otimes I_T\right)(\theta - \hat{\theta})X + \text{terms independent of } \theta$$

where $\hat{\theta} = \left[X'\left(\Sigma^{-1} \otimes I_T\right)X\right]^{-1} X'\left(\Sigma^{-1} \otimes I_T\right)y$.

[25] These convergence tools are implemented in Matlab Econometrics Toolbox, written by James P. LeSage (see www.spatial-econometrics.com).

The conditional posterior for θ is thus proportional to the terms in the exponents,

$$(\theta - \hat{\theta})' X' \left(\Sigma^{-1} \otimes I_T\right) (\theta - \hat{\theta}) X + (\theta - A_0\bar{\theta})'\Delta^{-1}(\theta - A_0\bar{\theta}).$$

This expression is similar to the standard multivariate distribution with an informative prior about θ (see e.g. Koop et al. (2007) p. 108-110). It can be rewritten as

$$(\hat{\theta} - A_0\bar{\theta})' X' \left(\Sigma^{-1} \otimes I_T\right) X V_\theta \Delta^{-1}(\hat{\theta} - A_0\bar{\theta}) + (\theta - m_\theta)' V_\theta^{-1}(\theta - m_\theta)$$

where

$$\begin{aligned} V_\theta &= \left(X' \left(\Sigma^{-1} \otimes I_T\right) X + \Delta^{-1}\right)^{-1} \\ m_\theta &= V_\theta \left(X' \left(\Sigma^{-1} \otimes I_T\right) y + \Delta^{-1} A_0 \bar{\theta}\right) \end{aligned}$$

θ only enters through the term $(\theta - m_\theta)' V_\theta^{-1}(\theta - m_\theta)$, we can thus write

$$p(\theta | D, \bar{\theta}, \Delta, \Sigma) \propto \exp\left\{-\frac{1}{2}(\theta - m_\theta)' V_\theta^{-1}(\theta - m_\theta)\right\}.$$

This is the kernel of a normal density and therefore $\theta | D, \bar{\theta}, \Delta, \Sigma \sim \mathcal{N}(m_\theta, V_\theta) \times \rho \in (-1, 1)$. It is easy to sample from the posterior by sampling from $\mathcal{N}(m_\theta, V_\theta)$ and rejecting the draws where $|\rho| > 1$.

When the equity premium constraint on θ is entertained, the prior on θ is defined as

$$p(\theta | \bar{\theta}, \Delta) \propto |\Delta|^{-N/2} \exp\left\{-\frac{1}{2}(\theta - A_0\bar{\theta})'\Delta^{-1}(\theta - A_0\bar{\theta})\right\}, \ \rho \in (-1, 1), \ \mu, \beta \in E_\theta$$

where E_θ is the set where return forecasts are all positive

$$E_\theta = \{\mu_i + \beta_i x_{i,t} \geq 0, i = 1, \ldots, N, t = 1, \ldots, T\}.$$

The posterior is therefore $\theta | D, \bar{\theta}, \Delta, \Sigma \sim \mathcal{N}(m_\theta, V_\theta) \times \rho \in (-1, 1) \times \mu, \beta \in E_\theta$. The equity premium restriction appears to be much more severe than the restriction that $|\rho|$ cannot be greater than one. It is therefore quite inefficient to rely on an accept-reject procedure to sample from the posterior. Griffiths (2003) shows that it is preferable to use a Metropolis-

Hastings algorithm in that case. The algorithm generates candidates value θ^* that are accepted with probability

$$r = \min\left[\frac{p(\theta^*|D,\bar{\theta},\Delta,\Sigma)}{p(\theta^{(s-1)}|D,\bar{\theta},\Delta,\Sigma)},1\right] \times \rho \in (-1,1) \times \mu,\beta \in E_\theta$$

where p is the density of the multivariate normal distribution $\mathcal{N}(m_\theta, V_\theta)$ and $\theta^{(s-1)}$ is a draw from a previous step of the algorithm. Initialize the algorithm with a $\theta^{(0)}$ obtained with non-linear least squares. Following Griffiths (2003), I use a proposal density identical to the true density, but with a covariance matrix multiplied by a scalar $c < 1$, so that the candidates value are accepted 40%-50% of the times.

2. Conditional posterior for Σ^{-1}

Let $S = (y - X\theta)(y - X\theta)'$. This is the kernel of the residual covariance matrix given θ. I derive first $p(\Sigma|D,\theta)$ by using the fact that the Jeffreys prior for $|\Sigma|$ is $|\Sigma|^{-(NG+1)/2}$ (see Zellner (1971), Chapter 8), and then use the factor structure given by Equations (4.4) and (4.5) to compute S. Observe that

$$\begin{aligned} p(\Sigma|D,\theta) &\propto p(D|\theta,\Sigma^{-1})p(\Sigma) \\ &\propto |\Sigma|^{T/2}\exp\left[-\frac{1}{2}\mathrm{tr}(S\Sigma^{-1})\right] \times |\Sigma|^{-(NG+1)/2}. \end{aligned}$$

One can recognize the kernel expression of Wishart distribution with scale matrix S and T degrees of freedom, and thus $\Sigma^{-1} \sim \text{Wishart}(S^{-1}, T)$.

In order to compute S, define $\bar{u}_t = \sum_{i=1}^N u_{i,t}$, $\bar{v}_t = \sum_{i=1}^N v_{i,t}$. First, obtain estimates of δ_i^u and δ_i^v by projecting the residuals on the common factors $\bar{u}_t$ and $\bar{v}_t$. Second, obtain estimates of γ, σ^x and σ^y from the factors and of γ_i, σ_i^x and σ_i^y from the residuals. Third, define

$$\delta = \begin{pmatrix} \delta_1^u \\ \delta_1^v \\ \vdots \\ \delta_N^u \\ \delta_N^v \end{pmatrix}, \quad V = \begin{pmatrix} (\sigma^x)^2 \\ (\sigma^y)^2 \\ \vdots \\ (\sigma^x)^2 \\ (\sigma^y)^2 \end{pmatrix}, \quad H = \begin{pmatrix} 1 & \gamma & \dots & 1 & \gamma \\ \gamma & 1 & & \gamma & 1 \\ \vdots & & \ddots & & \vdots \\ 1 & \gamma & & 1 & \gamma \\ \gamma & 1 & \dots & \gamma & 1 \end{pmatrix}$$

$$\Lambda = \begin{pmatrix} (\sigma_1^x)^2 & \gamma_1\sigma_1^x\sigma_1^y & \dots & & 0 \\ \gamma_1\sigma_1^x\sigma_1^y & (\sigma_1^y)^2 & & & \\ \vdots & & \ddots & & \vdots \\ & & & (\sigma_N^x)^2 & \gamma_N\sigma_N^x\sigma_N^y \\ 0 & & \dots & \gamma_N\sigma_N^x\sigma_N^y & (\sigma_N^y)^2 \end{pmatrix}$$

Then $S = T(V'V \cdot H \cdot (\delta'\delta) + \Lambda)$.

3. Conditional Posterior for $\bar{\theta}$

Using the fact that $\theta = A_0\bar{\theta} + \eta$, with a diffuse prior, it can be verified (see, e.g., Smith (1973)) that the conditional posterior is normal with mean

$$m_{\bar{\theta}} = V_{\bar{\theta}}\left(A_0'\Delta^{-1}\theta\right)$$

and covariance

$$V_{\bar{\theta}} = \left(A_0'\Delta^{-1}A_0\right)^{-1}.$$

4. Conditional Posterior for Δ^{-1} and T^{-1}

From Equation (23), and using the definition of Δ^{-1}, the conditional posterior density for the dispersion parameters is given by:

$$\begin{aligned} p(\Delta^{-1}|\theta,\bar{\theta},D) &\propto p(\theta|\bar{\theta},\Delta^{-1})p(\Delta^{-1}) \propto \prod_{i=1}^{N} p(\theta_i|\bar{\theta},\mathrm{T}^{-1})p(\mathrm{T}^{-1}) \\ &\propto |\mathrm{T}^{-1}|^{N/2}\exp\left\{-\frac{1}{2}\sum_{i=1}^{N}(\theta_i-\bar{\theta})'\mathrm{T}^{-1}(\theta_i-\bar{\theta})\right\} \times |\mathrm{T}^{-1}|^{-(G+1)/2} \\ &\propto |\mathrm{T}^{-1}|^{(N-G)/2}\exp\left\{-\frac{1}{2}\mathrm{tr}\left[\sum_{i=1}^{N}(\theta_i-\bar{\theta})'(\theta_i-\bar{\theta})\right]\mathrm{T}^{-1}\right\} \end{aligned}$$

Therefore $\mathrm{T}^{-1} \sim \text{Wishart}\left(\left[\sum_{i=1}^{N}(\theta_i-\bar{\theta})'(\theta_i-\bar{\theta})\right], N\right)$.

Inference when the number of observations is unequal

I just described the inference procedure under the assumption of a balanced panel of countries. Addressing the unbalanced nature of the data is relatively straightforward.

Out of simplicity I will consider the simpler case of a 2-country panel, where data is observed over the periods $t = 0, \dots T$ for country 1 and $t = t_0, \dots, T$ for country 2:

$$y_1 = X_1\theta_1 + \epsilon_1 \tag{24}$$

$$y_2 = X_2\theta_2 + \epsilon_2 \tag{25}$$

where $y_1 = (r_{1,1}, r_{1,2}, \dots r_{1,T}, x_{1,1}, x_{1,2}, \dots x_{1,T})'$, $y_2 = (r_{2,t_0+1}, r_{2,2}, \dots r_{2,T}, x_{2,t_0+1}, x_{2,2}, \dots x_{2,T})'$, $X_1 = \mathrm{diag}([((1, x_{1,0})', (1, x_{1,1})', \dots, (1, x_{1,T-1})')], [((1, x_{1,0})', (1, x_{1,1})', \dots, (1, x_{1,T-1})')])$, $X_2 = ((1, x_{2,t_0})', (2, x_{1,t_0+1})', \dots, (1, x_{2,T-1})')'$, $\epsilon_1 = (u_{1,1}, u_{1,2}, \dots u_{1,T}, v_{1,1}, v_{1,2}, \dots v_{1,T})'$ and $\epsilon_2 = (u_{2,t_0+1}, u_{2,2}, \dots u_{2,T}, v_{2,t_0+1}, v_{2,2}, \dots v_{2,T})'$. For a balanced panel, the error covariance matrix is given by $\Omega = \Sigma \otimes I_T$, while in our case it is given by

$$\Omega = \begin{bmatrix} \sigma_{11} \otimes I_{t0-1} & 0 & 0 \\ 0 & \sigma_{11} \otimes I_{T-t0+1} & \sigma_{12} \otimes I_{T-t0+1} \\ 0 & \sigma_{12} \otimes I_{T-t0+1} & \sigma_{22} \otimes I_{T-t0+1} \end{bmatrix} \tag{26}$$

where

$$\Sigma = \begin{pmatrix} \sigma_{11} & \sigma_{12} \\ \sigma_{12} & \sigma_{22} \end{pmatrix}$$

It is thus straightforward to handle unbalanced data within the SUR model. The only caveat is that the innovations $\epsilon_i = y_i - X_i\theta_i$ have the form

$$\epsilon = \begin{bmatrix} \epsilon_1^0 & 0 \\ \epsilon_1^1 & \epsilon_2^1 \end{bmatrix} \tag{27}$$

where the first line of the partition corresponds to the observations at times $t < t_0$. In order to draw values of Σ, it is thus necessary to compute the scale matrix S on the overlapping values, i.e. to ignore the information in e_0^i.

Bibliography

Andres, Christian, André Betzer, Marc Goergen, and Luc Renneboog, 2009, Dividend policy of German firms, *Journal of Empirical Finance* 16, 175–187.

Ang, Andrew, and Geert Bekaert, 2007, Stock Return Predictability: Is it There?, *Review of Financial Studies* 20, 651–707.

Avramov, Doron, 2002, Stock return predictability and model uncertainty, *Journal of Financial Economics* 64, 423–458.

Bansal, Ravi, Dana Kiku, and Amir Yaron, 2012, An empirical evaluation of the long-run risks model for asset prices, *Critical Finance Review* 1, 183–221.

Bansal, Ravi, and Amir Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481–1509.

Barberis, Nicholas, 2000, Investing for the long run when returns are predictable, *Journal of Finance* 55, 225–264.

Bec, Frédérique, and Christian Gollier, 2009, Assets Returns Volatility and Investment Horizon : The French Case, *CESifo Working Paper* 2622.

Boudoukh, Jacob, Matthew Richardson, and Robert F Whitelaw, 1997, Nonlinearities the Relation Premium and Between the the Equity Risk Term Structure, *Management Science* 43, 371–385.

Cales, Ludovic, Eric Jondeau, and Michael Rockinger, 2013, Long-Term Portfolio Management with a Structural Macroeconomic Model, *Swiss Finance Institute Research Paper* 13-45.

Campbell, John Y., 2003, Consumption-based asset pricing, in G.M. Constantinides, M. Harris, and R. Stulz, eds., *Handbook of the Economics of Finance*, 803–888 (North-

Holland, Armsterdam).

Campbell, John Y., and John H. Cochrane, 1999, By force of habit: A consumption-based explanation of aggregate stock market behavior, *Journal of Political Economy* 107, 205–251.

Campbell, John Y., and Robert J. Shiller, 1988, Stock Prices, Earnings, and Expected Dividends, *Journal of Finance* 43, 661–676.

Campbell, John Y., and Robert J. Shiller, 1991, Yield spreads and interest rate movements: A bird's eye view, *The Review of Economic* 58, 495–514.

Campbell, John Y., and Samuel B. Thompson, 2008, Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average?, *Review of Financial Studies* 21, 1509–1531.

Campbell, John Y., and Luis M. Viceira, 2002, *Strategic Asset Allocation*, volume 2002 (Oxford University Press, Oxford, UK).

Campbell, John Y., and Luis M. Viceira, 2004, *Long-horizon mean-variance analysis: A user guide*, Manuscript, Harvard University, Cambridge, MA.

Campbell, John Y., and Luis M Viceira, 2005, The term structure of the risk-return tradeoff, *Financial Analysts Journal* 61, 34–44.

Canova, Fabio, 2004, Testing for Convergence Clubs in Income Per Capita: A Predictive Density Approach, *International Economic Review* 45, 49–77.

Chen, Long, 2009, On the reversal of return and dividend growth predictability: A tale of two periods, *Journal of Financial Economics* 92, 128–151.

Chen, Long, Zhi Da, and Richard Priestley, 2012, Dividend smoothing and predictability, *Management Science* 58, 1834–1853.

Chib, Siddhartha, and Edward Greenberg, 1995a, Hierarchical analysis of SUR models with extensions to correlated serial errors and time-varying parameter models, *Journal of Econometrics* 68, 339–360.

Chib, Siddhartha, and Edward Greenberg, 1995b, Understanding the metropolis-hastings algorithm, *The American Statistician* 49, 327–335.

Clark, Todd E., and Kenneth D. West, 2007, Approximately normal tests for equal predictive accuracy in nested models, *Journal of Econometrics* 138, 291–311.

Cochrane, John H., 2008, The Dog That Did Not Bark: A Defense of Return Predictability, *Review of Financial Studies* 21, 1533–1575.

Cochrane, John H., 2011, Presidential address: discount rates, *Journal of Finance* 66, 1047–1108.

Cremers, K.J. Martijn, 2002, Stock return predictability: A Bayesian model selection perspective, *Review of Financial Studies* 15, 1223–1249.

Dangl, Thomas, and Michael Halling, 2012, Predictive regressions with time-varying coefficients, *Journal of Financial Economics* 106, 157–181.

de Finetti, Bruno, 1964, Foresight: its logical laws in subjective sources, in Jr. H.E. Kyburg, and H.E. Smokler, eds., *Studies in Subjective Probability*, 93–158 (Wiley, New York).

Diebold, Francis X., and Robert S. Mariano, 1995, Comparing Predictive Accuracy, *Journal of Business and Economic Statistics* 13, 253–263.

Dimson, Elroy, Paul Marsh, and Mike Staunton, 2008, The worldwide equity premium: a smaller puzzle, in Rajnish Mehra, ed., *Handbook Of The Equity Risk Premium*, 467–514 (Elsevier).

Diris, Bart F., 2011, Model uncertainty for long-term investors.

Driesprong, Gerben, Ben Jacobsen, and Benjamin Maat, 2008, Striking oil: Another puzzle?, *Journal of Financial Economics* 89, 307–327.

Engsted, Tom, and Thomas Q. Pedersen, 2010, The dividend-price ratio does predict dividend growth: International evidence, *Journal of Empirical Finance* 17, 585–605.

Fama, Eugene F., 1991, Efficient Capital Markets: II, *Journal of Finance* 46, 1575–1617.

Fama, Eugene F., and Kenneth R. French, 1988, Dividend yields and expected stock returns, *Journal of Financial Economics* 22, 3–25.

Fama, Eugene F., and Kenneth R. French, 1989, Business conditions and expected returns on stocks and bonds, *Journal of financial economics* 25, 23–49.

Fama, Eugene F., and Kenneth R. French, 2001, Disappearing dividends: changing firm characteristics or lower propensity to pay?, *Journal of Financial Economics* 60, 3–43.

Fama, Eugene F., and G. William Schwert, 1977, Asset returns and inflation, *Journal of Financial Economics* 5, 115–146.

Favero, Carlo A., and Andrea Tamoni, 2010, Demographics and US Stock Market Fluctuations, *CESifo Economic Studies* 57, 25–43.

Gabaix, Xavier, 2008, Variable rare disasters: A tractable theory of ten puzzles in macro-finance, *The American Economic Review* 98.

Geweke, John, 1992, Evaluating the Accuracy of Sampling-Based Approaches to the Calculation of Posterior Moments, in J. M. Bernardo, J. O. Berger, A. P. Dawid, and A. F. M. Smith, eds., *Bayesian Statistics, Volume 4*, 169–193 (Oxford University Press).

Griffiths, William E., 2003, Bayesian Inference in the Seemingly Unrelated Regressions Model, in David Giles, ed., *Computer-Aided Econometrics* (CRC Press).

Henkel, Sam James, J. Spencer Martin, and Federico Nardari, 2011, Time-varying short-horizon predictability, *Journal of Financial Economics* 99, 560–580.

Hjalmarsson, Erik, 2010, Predicting global stock returns, *Journal of Financial and Quantitative Analysis* 45, 49–80.

Hoevenaars, Roy P. P. M., Roderick D. J. Molenaar, Peter C. Schotman, and Tom B. M. Steenkamp, 2014, Strategic Asset Allocation for Long-Term Investors: Parameter Uncertainty and Prior Information, *Journal of Applied Econometrics* 29, 353–376.

Hoogstrate, Andre, J. Franz C. Palm, and Gerard A. Pfann, 2000, Pooling in dynamic panel-data models: An application to forecasting GDP growth rates, *Journal of Business & Economic Statistics* 18, 274–283.

Hsiao, Cheng, and M. Hashem Pesaran, 2004, Random coefficient panel data models, *CESifo working papers* 1233.

Hsiao, Cheng, M. Hashem Pesaran, and A. Kamil Tahmiscioglu, 1998, Bayes estimation of short-run coefficients in dynamic panel data models, in C. Hsiao, K. Lahiri, L.-F. Lee, and M.H. Pesaran, eds., *Analysis of Panels and Limited Dependent Variables: A Volume in Honour of G. S. Maddala*, number 96, chapter 11, 268–296 (Cambridge University Press, Cambridge).

Jeffreys, H, 1961, *The theory of probability* (Oxford University Press, Clarenden).

Johannes, Michael, Arthur Korteweg, and Nicholas Polson, 2014, Sequential Learning, Predictability, and Optimal Portfolio Returns, *Journal of Finance* forthcoming.

Johannes, Michael, and Nicholas Polson, 2003, MCMC methods for continuous-time financial econometrics, in Yacine Ait-Sahalia, and Lars Hansen, eds., *Handbook of Financial Econometrics* (Elsevier, North-Holland).

Jondeau, Eric, and Michael Rockinger, 2010, Portfolio Allocation for European Markets with Predictability and Parameter Uncertainty, *Swiss Finance Institute Research Paper* 10.

Jones, Christopher S., and Jay Shanken, 2005, Mutual fund performance with learning across funds, *Journal of Financial Economics* 78, 507–552.

Kandel, Shmuel, and Robert F. Stambaugh, 1996, On the Predictability of Stock Returns: An Asset-Allocation Perspective, *Journal of Finance* 51, 385–424.

Keim, Donald B., and Robert F. Stambaugh, 1986, Predicting returns in the stock and bond markets, *Journal of Financial Economics* 17, 357–390.

Kelly, Bryan, and Seth Pruitt, 2012, The Three-Pass Regression Filter: A New Approach to Forecasting Using Many Predictors.

Koijen, Ralph S. J., and Jules H. Van Binsbergen, 2010, Predictive Regressions: A Present-Value Approach, *Journal of Finance* 65, 1439–1471.

Koijen, Ralph S.J., and Stijn Van Nieuwerburgh, 2011, Predictability of Returns and Cash Flows, *Annual Review of Financial Economics* 3, 467–491.

Koop, Gary, Dale J. Poirier, and Justin L. Tobias, 2007, *Bayesian econometric methods* (Cambridge University Press).

Kothari, Smitu P., and Jay Shanken, 1997, Book-to-market, dividend yield, and expected market returns: A time-series analysis, *Journal of Financial Economics* 44, 169–203.

Lettau, M, and S Ludvigson, 2001, Consumption, aggregate wealth, and expected stock returns, *Journal of Finance* 56, 815–849.

Lettau, M., and S. Van Nieuwerburgh, 2008, Reconciling the Return Predictability Evidence, *Review of Financial Studies* 21, 1607–1652.

Lindley, Dennis V., and Adrian F.M. Smith, 1972, Bayes Estimates for the Linear Model, *Journal of the Royal Statistical Society* 34, 1–41.

Ludvigson, Sydney C., and Serena Ng, 2007, The empirical riskreturn relation: A factor analysis approach, *Journal of Financial Economics* 83, 171–222.

Maio, Paulo F., and P Santa-Clara, 2012, Dividend yields, dividend growth, and return predictability in the cross-section of stocks, *Journal of Financial and Quantitative Analysis forthcoming* .

Menzly, Lior, Tano Santos, and Pietro Veronesi, 2004, Understanding Predictability, *Journal of Political Economy* 112, 1–47.

Merton, Robert C., 1980, On estimating the expected return on the market, *Journal of Financial Economics* 8, 323–361.

Mittnik, S, 1990, Macroeconomic forecasting using pooled international data, *Journal of Business & Economic Statistics* 8, 205–208.

Neely, Christopher J., David E. Rapach, Jun Tu, and Guofu Zhou, 2014, Forecasting the equity risk premium: the role of technical indicators, *Management Science* 60, 1772–1791.

Pastor, Lubos, and Robert F. Stambaugh, 2009, Predictive systems: Living with imperfect predictors, *Journal of Finance* LXIV.

Pastor, Lubos, and Robert F. Stambaugh, 2012, Are Stocks Really Less Volatile in the Long Run?, *Journal of Finance* 67, 431–478.

Paye, Bradley S., and Allan G. Timmermann, 2006, Instability of return prediction models, *Journal of Empirical Finance* 13, 274–315.

Pesaran, M. Hashem, and Allan G. Timmermann, 2002, Market timing and return prediction under model instability, *Journal of Empirical Finance* 9, 495–510.

Pettenuzzo, Davide, and Allan G. Timmermann, 2011, Predictability of stock returns and asset allocation under structural breaks, *Journal of Econometrics* 164, 60–78.

Pettenuzzo, Davide, Allan G. Timmermann, and Rossen Valkanov, 2013, Forecasting Stock Returns under Economic Constraints, *Journal of Financial Economics (forthcoming)* .

Raftery, Adrian E., and Steven Lewis, 1992a, How many iterations in the Gibbs sampler, *Bayesian statistics* 4, 763–773.

Raftery, Adrian E., and Steven M. Lewis, 1992b, One Long Run with Diagnostics: Implementation Strategies for Markov Chain Monte Carlo, *Statistical Science* 7, 493–497.

Raftery, Adrian E., and Steven M. Lewis, 1995, The Number of Iterations, Convergence Diagnostics and Generic Metropolis Algorithms, in W.R. Gilks, and D.J. Spiegelhalter, eds., *Practical Markov Chain Monte Carlo*, 115–130 (Chapman and Hall).

Rapach, David E., 2006, Structural Breaks and Predictive Regression Models of Aggregate U.S. Stock Returns, *Journal of Financial Econometrics* 4, 238–274.

Rapach, David E., Jack K. Strauss, and Guofu Zhou, 2009, Out-of-Sample Equity Premium Prediction: Combination Forecasts and Links to the Real Economy, *Review of Financial Studies* 23, 821–862.

Rapach, David E., Jack K. Strauss, and Guofu Zhou, 2013, International Stock Return Predictability: What Is the Role of the United States?, *Journal of Finance* 68, 1633–1662.

Rapach, David E., Mark E. Wohar, and Jesper Rangvid, 2005, Macro variables and international stock return predictability, *International Journal of Forecasting* 21, 137–166.

Rapach, David E., and Guofu Zhou, 2012, Forecasting Stock Returns, in Graham Elliott, and Allan Timmermann, eds., *Handbook of Economic Forecasting, Volume 2*, volume 2, chapter Forecastin (North Holland).

Renneboog, Luc, and Grzegorz Trojanowski, 2007, Control structures and payout policy, *Managerial Finance* 33, 43–64.

Ross, Steven A., 2005, *Neoclassical finance* (Princeton University Press, Princeton, NJ).

Rozeff, Michael, 1984, Dividend yields are equity risk premiums, *Journal of Portfolio management* 68–75.

Schrimpf, Andreas, 2010, International stock return predictability under model uncertainty, *Journal of International Money and Finance* 29, 1256–1282.

Schwert, G. William, 2003, Anomalies and market efficiency, in George Constantinides, Milton Harris, and René Stulz, eds., *Handbook of the Economics of Finance*, north-holl edition, chapter 15, 937–972.

Siegel, Jeremy J., 1998, *Stocks for the long run*, mcgraw hil edition (New York).

Smith, Adrian F. M., 1973, A general Bayesian linear model, *Journal of the Royal Statistical Society. Series B (. . .* 35, 67–75.

Stambaugh, Robert F., 1999, Predictive Regressions, *Journal of Financial Economics* 54, 375–421.

Von Eije, Henk, and William L. Megginson, 2008, Dividends and share repurchases in the European Union, *Journal of Financial Economics* 89, 347–374.

Wachter, Jessica A., and Missaka Warusawitharana, 2009, Predictable Returns and Asset Allocation: Should a Skeptical Investor Time the Market?, *Journal of Econometrics* 148, 162–178.

Welch, Ivo, and Amit Goyal, 2008, A Comprehensive Look at The Empirical Performance of Equity Premium Prediction, *Review of Financial Studies* 21, 1455–1508.

West, Kenneth D., 2006, Forecast Evaluation, in G. Elliott, C.W.J. Granger, and A. Timmermann, eds., *Handbook of Economic Forecasting, Volume 1*, volume 1, chapter 3, 99–134 (Elsevier, Amsterdam).

Zellner, Arnold, 1962, An efficient method of estimating seemingly unrelated regressions and tests for aggregation bias, *Journal of the American Statistical Association* 57, 348–368.

Zellner, Arnold, 1971, *An introduction to Bayesian inference in econometrics* (Wiley, New York).

Zellner, Arnold, and Chansik Hong, 1989, Forecasting international growth rates using Bayesian shrinkage and other procedures, *Journal of Econometrics* 40, 183–202.

Zhou, Guofu, 2010, How much stock return predictability can we expect from an asset pricing model?, *Economics Letters* 108, 184–186.

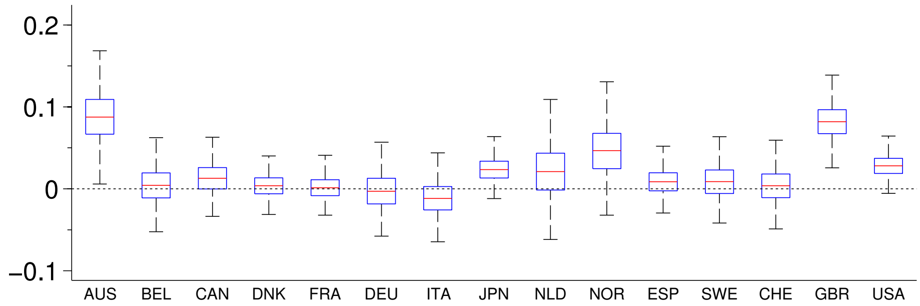

(a) No Cross-sectional Learning

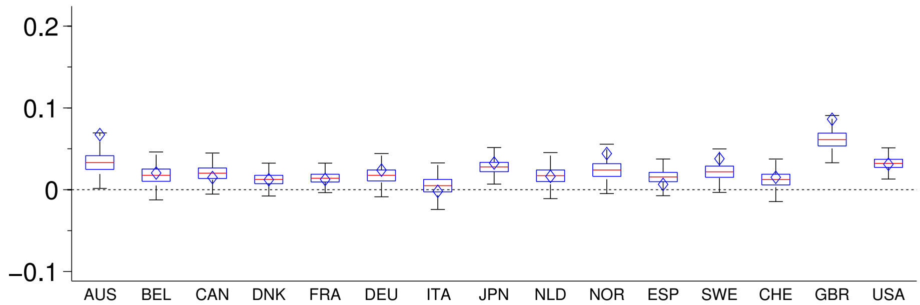

(b) Cross-sectional Learning

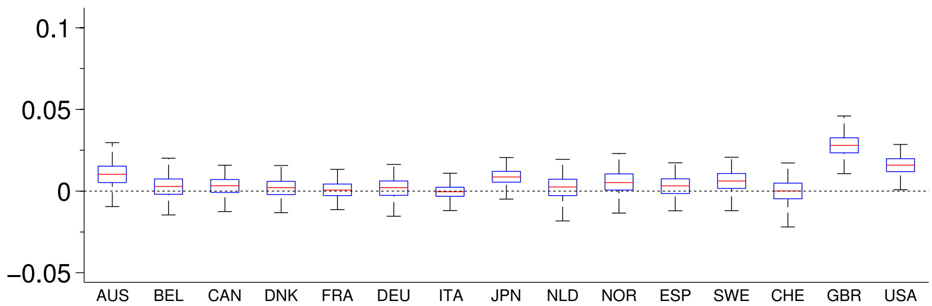

(c) Cross-sectional Learning and Equity Premium Constraint

Figure 4.1: Stock Return, Posterior Distribution for the Dividend-Price Ratio Coefficient

This figure depicts the posterior distributions of the slope coefficient β in a regression of excess returns on the log-dividend-price ratio, for each of the fifteen countries. For each country the red bar represents the median, the box corresponds to the first and third quartiles, and the whiskers represent the 99% confidence band. The diamonds in Panel (b) indicate the frequentist SUR estimates.

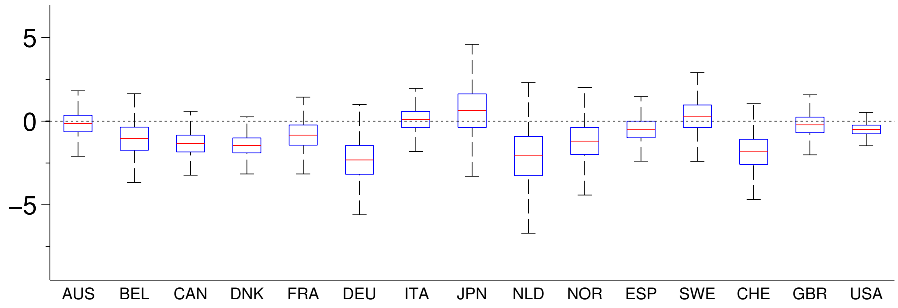

(a) No Cross-sectional Learning

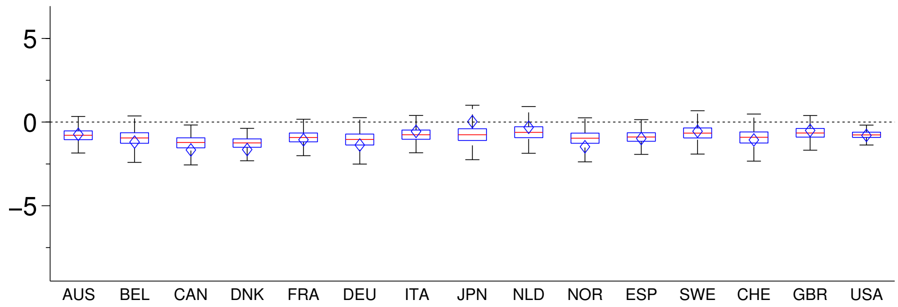

(b) Cross-sectional Learning

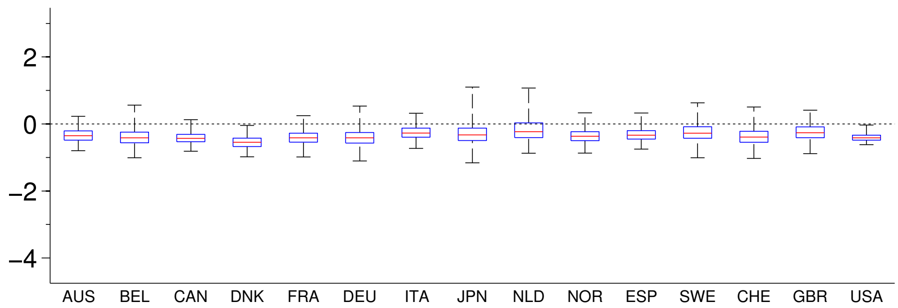

(c) Cross-sectional Learning and Equity Premium Constraint

Figure 4.2: Stock Return, Posterior Distribution for the Short-term Interest Rate Coefficient
This figure depicts the posterior distributions of the slope coefficient β in a regression of excess returns on the short-term interest rate, for each of the fifteen countries. For each country the red bar represents the median, the box corresponds to the first and third quartiles, and the whiskers represent the 99% confidence band. The diamonds in Panel (b) indicate the frequentist SUR estimates.

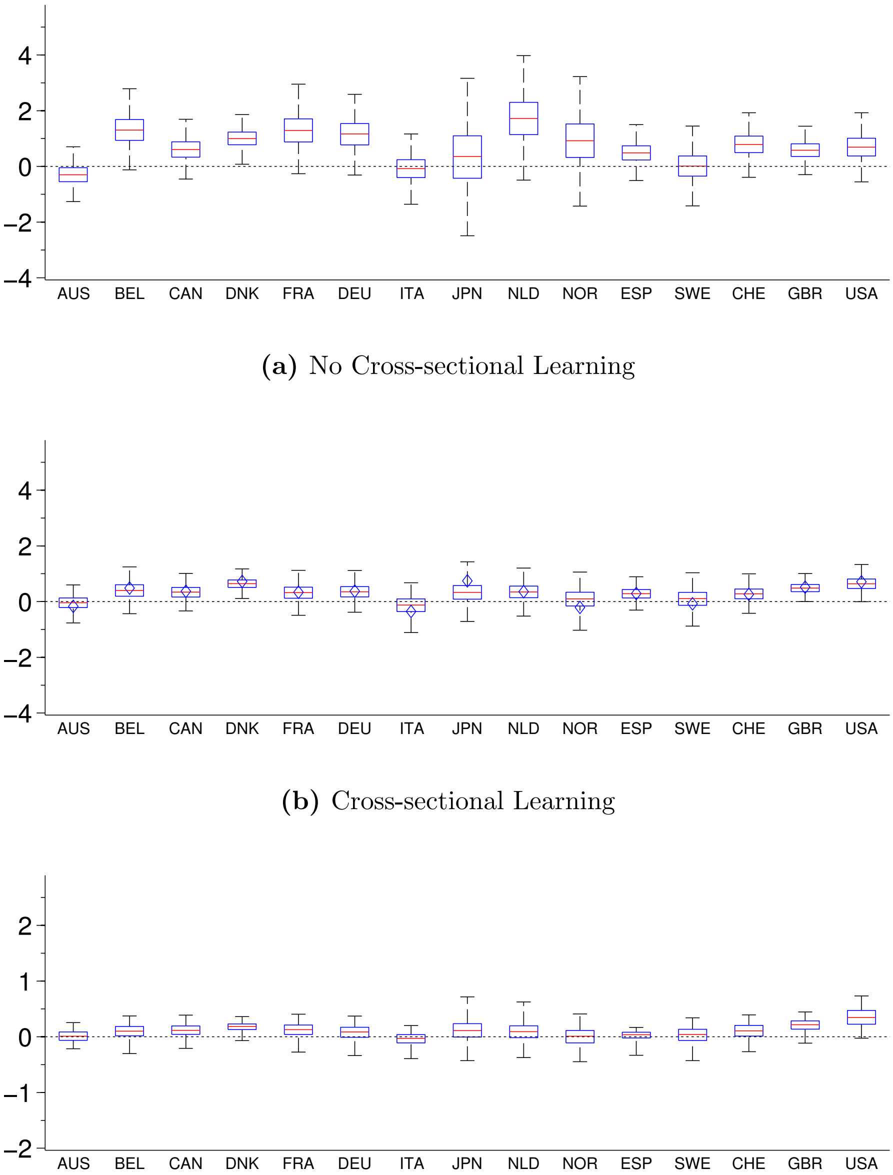

(a) No Cross-sectional Learning

(b) Cross-sectional Learning

(c) Cross-sectional Learning and Equity Premium Constraint

Figure 4.3: Stock Return, Posterior Distribution for the Term Spread Coefficient

This figure depicts the posterior distributions of the slope coefficient β in a regression of excess returns on the term spread, for each of the fifteen countries. For each country the red bar represents the median, the box corresponds to the first and third quartiles, and the whiskers represent the 99% confidence band. The diamonds in Panel (b) indicate the frequentist SUR estimates.

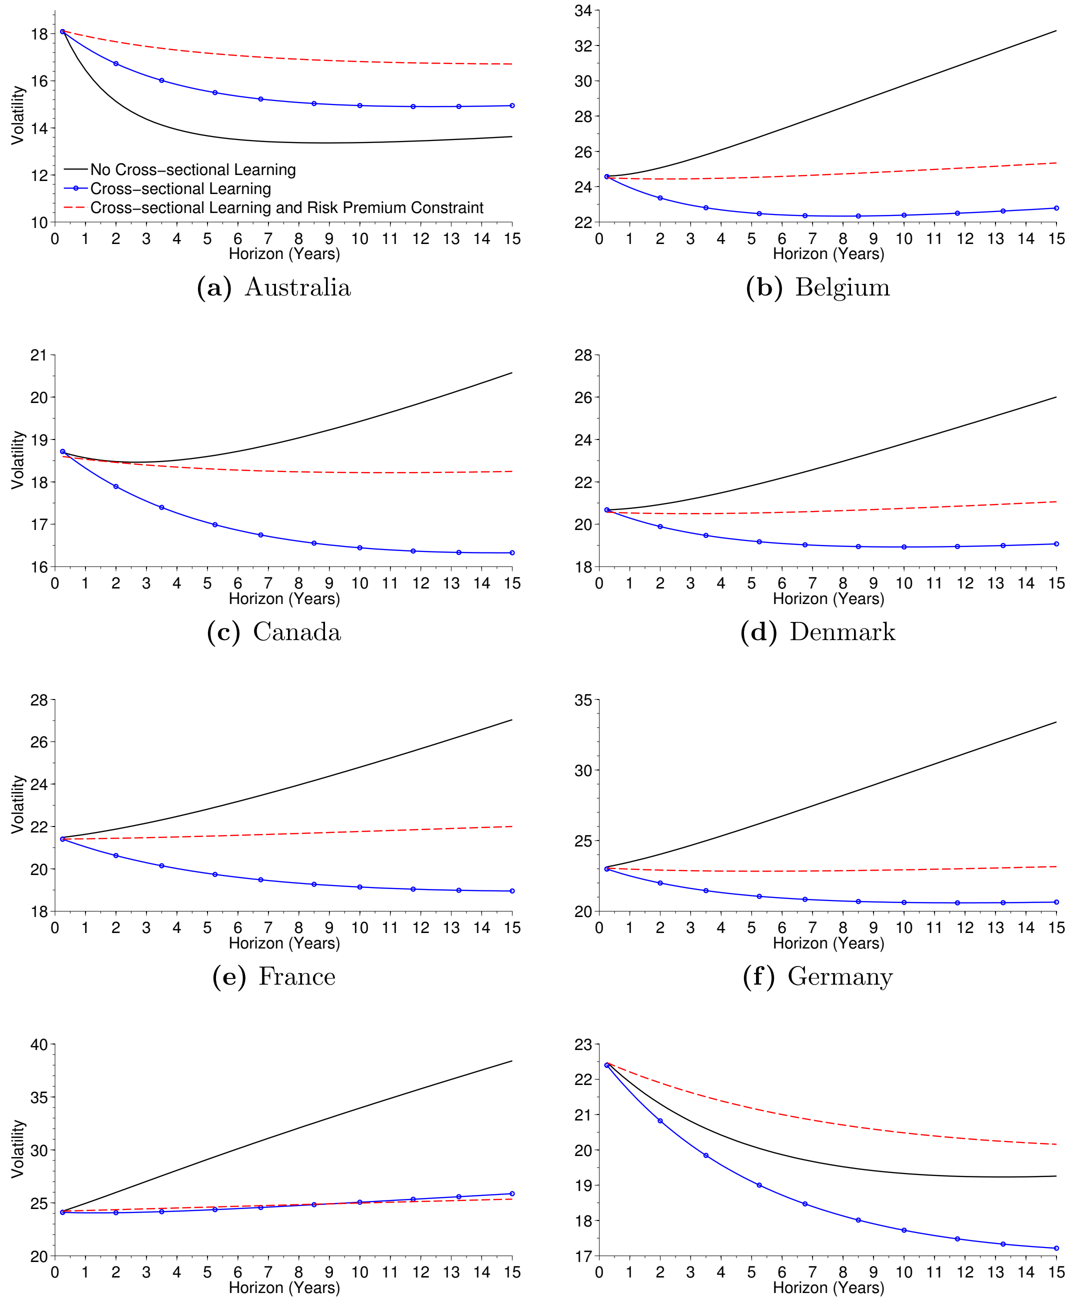

(a) Australia **(b)** Belgium **(c)** Canada **(d)** Denmark **(e)** France **(f)** Germany **(g)** Italy **(h)** Japan

Figure 4.4: Term Structure of Risk

This figure shows the term structure of annualized predictive excess return volatilities for each country and for the three models. The straight lines use only country-level information. The dotted lines correspond to the exchangeable prior, i.e. the model partially pools cross-country coefficients. The dashed lines feature cross-sectional learning and constrain the posterior distribution of the equity premium to be positive.

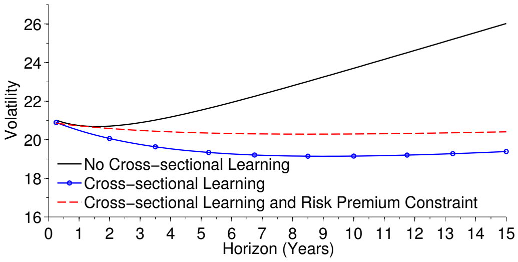

(a) The Netherlands

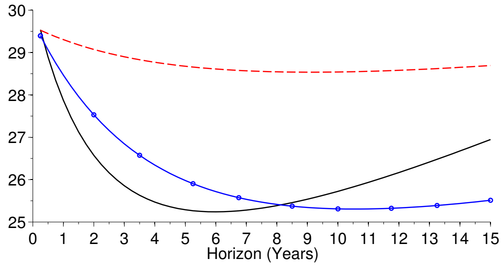

(b) Norway

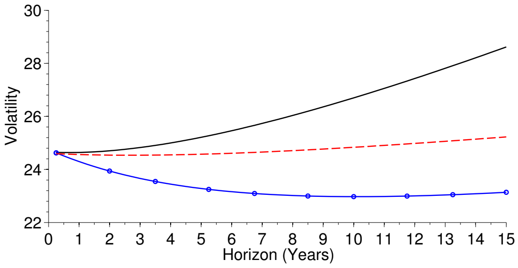

(c) Spain

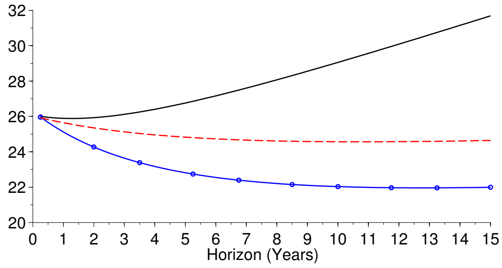

(d) Sweden

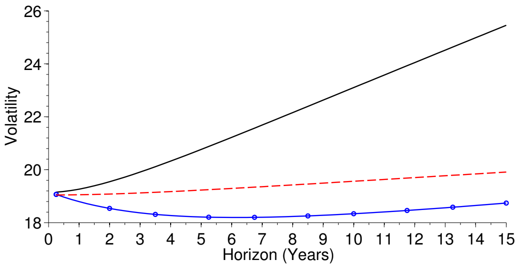

(e) Switzerland

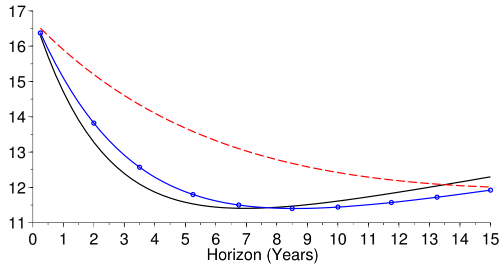

(f) The United Kingdom

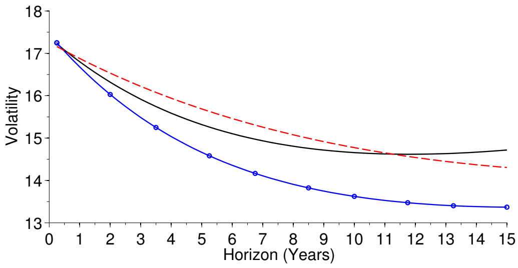

(g) The USA

Figure 4.5: Term Structure of Risk (Continued)
This figure shows the term structure of annualized predictive excess return volatilities for each country and for the three models. The straight lines use only country-level information. The dotted lines correspond to the exchangeable prior, i.e. the model partially pools cross-country coefficients. The dashed lines feature cross-sectional learning and constrain the posterior distribution of the equity premium to be positive.

Table I: Summary Statistics

	x_e	g	g^r	dp	dp^r	y	spr
Australia	4.539	-0.377	2.032	-3.233	-3.179	7.784	0.538
Q1-1971 - Q1-2013	(20.024)	(8.759)	(20.635)	(0.113)	(0.105)	(1.898)	(0.94)
Belgium	6.196	0.783	1.522	-3.095	-3.044	6.079	0.967
Q1-1971 - Q1-2013	(21.579)	(26.144)	(27.354)	(0.207)	(0.198)	(1.829)	(0.812)
Canada	4.649	-0.376	0.994	-3.606	-3.551	6.565	0.983
Q1-1971 - Q1-2013	(18.454)	(6.09)	(17.906)	(0.182)	(0.175)	(1.928)	(0.869)
Denmark	7.128	1.206	2.899	-3.815	-3.744	7.465	1.168
Q1-1972 - Q1-2013	(20.165)	(22.7)	(27.667)	(0.279)	(0.275)	(2.4)	(1.167)
France	5.435	0.45	2.912	-3.536	-3.484	6.52	0.827
Q1-1971 - Q1-2013	(21.001)	(12.642)	(23.915)	(0.269)	(0.259)	(1.897)	(0.663)
Germany	5.078	0.611	3.312	-3.658	-3.62	5.017	0.996
Q1-1971 - Q1-2013	(20.042)	(13.523)	(24.646)	(0.189)	(0.173)	(1.415)	(0.775)
Italy	2.047	-1.848	0.411	-3.579	-3.541	8.55	0.618
Q1-1971 - Q1-2013	(24.482)	(19.946)	(27.056)	(0.216)	(0.177)	(2.622)	(0.971)
Japan	5.564	0.273	2.73	-4.374	-4.344	2.558	1.809
Q1-1971 - Q1-2013	(19.925)	(10.412)	(21.486)	(0.28)	(0.281)	(1.15)	(0.392)
Netherlands	6.637	1.083	2.931	-3.4	-3.353	4.072	1.094
Q1-1986 - Q1-2013	(20.5)	(16.115)	(24.508)	(0.151)	(0.135)	(1.177)	(0.571)
Norway	9.013	3.61	6.938	-3.586	-3.524	7.574	-0.389
Q1-1979 - Q1-2013	(29.078)	(22.544)	(34.048)	(0.201)	(0.187)	(2.064)	(0.694)
Spain	7.375	-1.424	1.582	-3.095	-3.023	7.884	0.79
Q2-1978 - Q1-2013	(23.034)	(10.567)	(23.461)	(0.311)	(0.306)	(2.803)	(1.379)
Sweden	10.539	5.663	8.846	-3.606	-3.527	6.688	0.94
Q1-1971 - Q1-2013	(23.822)	(19.058)	(30.138)	(0.234)	(0.214)	(1.993)	(0.916)
Switzerland	6.419	5.131	7.283	-3.845	-3.801	3.315	0.542
Q1-1974 - Q1-2013	(17.194)	(21.916)	(27.56)	(0.178)	(0.168)	(1.357)	(0.855)
United Kingdom	6.434	-0.229	1.883	-3.282	-3.218	7.067	1.115
Q1-1971 - Q1-2013	(20.465)	(6.244)	(20.42)	(0.142)	(0.145)	(1.853)	(0.969)
United States	10.764	4.857	6.024	-3.514	-3.454	8.023	1.439
Q1-1953 - Q1-2013	(15.261)	(4.75)	(16.375)	(0.198)	(0.201)	(2.954)	(0.604)

This table reports the mean and standard deviations (in parenthesis) for the log excess stock return (x_e), log dividend growth (g), log dividend-price ratio (dp), the short-term interest rate (y) and the term spread (spr). The subscript on g^r and dp^r indicates that the variables are constructed with dividends reinvested. All variables (except dp, dp^r and spr) are expressed in percent per year. Excess returns and dividend growth are adjusted for Jensen's inequality by adding one-half of the sample variance.

Table II: Bayesian Estimates: Present-Value Model

		Dividends unreinvested			Dividends reinvested		
		Excess returns		Div. growth	Excess returns		Div. growth
Prior		Neutral	EP		Neutral	EP	
Country	Sample						
Australia	Q1-1971 - Q1-2013	0.033	0.010	-0.041	0.048	0.015	0.005
(Stdev)		(0.013)	(0.008)	(0.010)	(0.014)	(0.007)	(0.012)
$p < 0$		0.000	0.090	1.000	0.000	0.020	0.320
Belgium	Q1-1971 - Q1-2013	0.018	0.003	-0.045	0.034	0.008	0.000
(Stdev)		(0.011)	(0.007)	(0.013)	(0.011)	(0.007)	(0.011)
$p < 0$		0.060	0.330	1.000	0.000	0.130	0.490
Canada	Q1-1971 - Q1-2013	0.020	0.003	-0.038	0.028	0.006	0.000
(Stdev)		(0.010)	(0.006)	(0.006)	(0.010)	(0.005)	(0.010)
$p < 0$		0.020	0.290	1.000	0.000	0.150	0.490
Denmark	Q1-1972 - Q1-2013	0.013	0.002	-0.043	0.028	0.009	-0.001
(Stdev)		(0.008)	(0.006)	(0.010)	(0.008)	(0.005)	(0.009)
$p < 0$		0.050	0.370	1.000	0.000	0.050	0.550
France	Q1-1971 - Q1-2013	0.014	0.001	-0.038	0.024	0.005	-0.003
(Stdev)		(0.007)	(0.005)	(0.007)	(0.007)	(0.004)	(0.009)
$p < 0$		0.020	0.450	1.000	0.000	0.120	0.620
Germany	Q1-1971 - Q1-2013	0.018	0.002	-0.047	0.036	0.010	0.001
(Stdev)		(0.010)	(0.006)	(0.010)	(0.011)	(0.006)	(0.011)
$p < 0$		0.040	0.380	1.000	0.000	0.050	0.450
Italy	Q1-1971 - Q1-2013	0.005	0.000	-0.056	0.031	0.007	-0.001
(Stdev)		(0.011)	(0.004)	(0.012)	(0.011)	(0.006)	(0.011)
$p < 0$		0.330	0.540	1.000	0.010	0.110	0.540
Japan	Q1-1971 - Q1-2013	0.028	0.009	-0.021	0.034	0.010	0.004
(Stdev)		(0.009)	(0.005)	(0.007)	(0.009)	(0.005)	(0.008)
$p < 0$		0.000	0.050	1.000	0.000	0.020	0.330
Netherlands	Q1-1986 - Q1-2013	0.017	0.002	-0.047	0.037	0.012	0.002
(Stdev)		(0.011)	(0.007)	(0.011)	(0.011)	(0.006)	(0.011)
$p < 0$		0.060	0.380	1.000	0.000	0.030	0.440
Norway	Q1-1979 - Q1-2013	0.024	0.005	-0.042	0.043	0.013	0.006
(Stdev)		(0.012)	(0.007)	(0.011)	(0.012)	(0.007)	(0.011)
$p < 0$		0.020	0.220	1.000	0.000	0.020	0.290
Spain	Q2-1978 - Q1-2013	0.016	0.003	-0.038	0.022	0.005	-0.005
(Stdev)		(0.009)	(0.006)	(0.007)	(0.009)	(0.006)	(0.010)
$p < 0$		0.040	0.320	1.000	0.010	0.220	0.670
Sweden	Q1-1971 - Q1-2013	0.022	0.006	-0.045	0.044	0.013	0.004
(Stdev)		(0.010)	(0.007)	(0.010)	(0.011)	(0.006)	(0.011)
$p < 0$		0.010	0.180	1.000	0.000	0.010	0.350
Switzerland	Q1-1974 - Q1-2013	0.012	0.000	-0.044	0.030	0.008	-0.001
(Stdev)		(0.010)	(0.007)	(0.012)	(0.010)	(0.005)	(0.010)
$p < 0$		0.110	0.490	1.000	0.000	0.080	0.550
UK	Q1-1971 - Q1-2013	0.061	0.028	-0.012	0.060	0.022	0.014
(Stdev)		(0.012)	(0.007)	(0.008)	(0.012)	(0.006)	(0.012)
$p < 0$		0.000	0.000	0.940	0.000	0.000	0.100
USA	Q1-1953 - Q1-2013	0.032	0.016	-0.010	0.033	0.013	0.005
(Stdev)		(0.007)	(0.006)	(0.006)	(0.007)	(0.005)	(0.009)
$p < 0$		0.000	0.000	0.960	0.000	0.010	0.300

This table reports Bayesian estimates for the regression $z_{i,t+1} = \mu_i + \beta_i x_{i,t} + u_{i,t+1}$ where the left hand-side variable is either excess returns or dividend growth and $x_{i,t}$ is the log dividend-price ratio. Neutral (EP) indicates that the forecasts of the equity premium are unconstrained (constrained) to be positive. The right panel corresponds to regressions where dividends are assumed to be reinvested in the stock market. Means, standard deviations, and the probability that β_i is smaller than zero are reported.

Table III: Bayesian Estimates: Other Predictors

		Short-term Interest Rate		Term spread	
Prior		Neutral	EP	Neutral	EP
Country	Sample				
Australia	Q1-1971 - Q1-2013	-0.787	-0.338	-0.048	0.012
(Stdev)		(0.411)	(0.206)	(0.258)	(0.103)
$p < 0$		0.970	0.930	0.570	0.470
Belgium	Q1-1971 - Q1-2013	-0.957	-0.393	0.399	0.098
(Stdev)		(0.494)	(0.271)	(0.318)	(0.125)
$p < 0$		0.980	0.930	0.100	0.210
Canada	Q1-1971 - Q1-2013	-1.248	-0.417	0.337	0.112
(Stdev)		(0.451)	(0.171)	(0.257)	(0.117)
$p < 0$		1.000	0.980	0.090	0.160
Denmark	Q1-1972 - Q1-2013	-1.262	-0.546	0.645	0.178
(Stdev)		(0.377)	(0.188)	(0.205)	(0.079)
$p < 0$		1.000	1.000	0.000	0.020
France	Q1-1971 - Q1-2013	-0.921	-0.401	0.321	0.122
(Stdev)		(0.404)	(0.221)	(0.302)	(0.132)
$p < 0$		0.990	0.950	0.140	0.170
Germany	Q1-1971 - Q1-2013	-1.051	-0.410	0.354	0.076
(Stdev)		(0.516)	(0.263)	(0.283)	(0.131)
$p < 0$		0.980	0.950	0.100	0.270
Italy	Q1-1971 - Q1-2013	-0.752	-0.254	-0.138	-0.042
(Stdev)		(0.421)	(0.207)	(0.341)	(0.115)
$p < 0$		0.960	0.880	0.640	0.610
Japan	Q1-1971 - Q1-2013	-0.733	-0.285	0.331	0.120
(Stdev)		(0.579)	(0.354)	(0.382)	(0.208)
$p < 0$		0.900	0.840	0.180	0.260
Netherlands	Q1-1986 - Q1-2013	-0.591	-0.168	0.347	0.094
(Stdev)		(0.510)	(0.352)	(0.325)	(0.178)
$p < 0$		0.880	0.730	0.140	0.280
Norway	Q1-1979 - Q1-2013	-0.972	-0.355	0.080	-0.002
(Stdev)		(0.477)	(0.222)	(0.388)	(0.165)
$p < 0$		0.980	0.930	0.390	0.470
Spain	Q2-1978 - Q1-2013	-0.891	-0.310	0.282	0.021
(Stdev)		(0.387)	(0.202)	(0.227)	(0.087)
$p < 0$		0.990	0.920	0.110	0.330
Sweden	Q1-1971 - Q1-2013	-0.646	-0.248	0.098	0.028
(Stdev)		(0.468)	(0.281)	(0.358)	(0.150)
$p < 0$		0.910	0.820	0.380	0.390
Switzerland	Q1-1974 - Q1-2013	-0.923	-0.371	0.275	0.102
(Stdev)		(0.515)	(0.268)	(0.269)	(0.136)
$p < 0$		0.970	0.910	0.150	0.220
UK	Q1-1971 - Q1-2013	-0.640	-0.246	0.484	0.208
(Stdev)		(0.398)	(0.247)	(0.193)	(0.109)
$p < 0$		0.940	0.840	0.010	0.040
USA	Q1-1953 - Q1-2013	-0.762	-0.401	0.645	0.349
(Stdev)		(0.228)	(0.111)	(0.259)	(0.169)
$p < 0$		1.000	1.000	0.010	0.010

This table reports Bayesian estimates of β_i for the predictive regressions $r_{i,t+1} = \mu_i + \beta_i x_{i,t} + u_{i,t+1}$ where $r_{i,t+1}$ are quarterly excess return and $x_{i,t}$ is the short-term interest rate (left panel) or the term spread (right panel). Neutral (EP) indicates that the forecasts of the equity premium are unconstrained (constrained) to be positive. Means, standard deviations, and the probability that β_i is larger than zero are reported.

Table IV: Out-of-Sample Forecasting Results

Country	Without Equity Premium Constraints LS (1)	Bayes (2)	Δ (3)	With Equity Premium Constraints LS (4)	Bayes (5)	Δ (6)
Panel A: Log-Dividend Price Ratio						
Australia	−6.17	−0.05	6.12	1.56	−0.16	−1.72
Belgium	−1.85	−1.49	0.36	−1.1	−0.7	0.48
Canada	−6.22	−2.11	4.10*	−1.62	2.26**	3.84**
Denmark	−2.25	−3.58	−1.33	−0.75	1.20*	2.05*
France	−0.96	−1.94	−0.98	−0.9	−0.11	0.73
Germany	−1.36	−1.78	−0.42	−0.47	−0.19	0.15
Italy	−2.04	−0.76	1.27	0.70*	−1.1	−1.71
Japan	0.36	0.8	0.43	0.19	0.44	0.32
Sweden	−2.92	−3.21	−0.29	−2.31	−0.57	1.74*
UK	−4.05	1.91	5.96	3.33*	1.41	−1.9
USA	−4.24	−1.71	2.53**	−3.57	1.65*	5.13**
Average	−2.88	−1.27	1.61	−0.45	0.38	0.83
Panel B: Short-term Interest Rate						
Australia	−1.8	1.11	2.91	−0.88	1.30*	2.18
Belgium	0.17	0.28	0.11	0.06	0.09	0.03
Canada	−0.76	0.54	1.3	0.58	0.73	0.14
Denmark	1.71*	2.35*	0.65	1.7	2.12**	0.42
France	−0.78	0.56	1.34*	−0.22	0.42	0.64
Germany	1.84*	1.62*	−0.22	0.63	1.19	0.55
Italy	−2.6	−0.19	2.41	0.36	−0.69	−1.05
Japan	−3.62	−2.72	0.89	−3.21	0.07	3.28*
Sweden	−3.64	−0.32	3.32	−2.42	−0.23	2.19
UK	−2.18	−2.75	−0.57	−1.98	−1.95	0.03
USA	−4.2	−3.68	0.51	−2.46	−1.04	1.43
Average	−1.44	−0.29	1.15	−0.71	0.18	0.89
Panel C: Term Spread						
Australia	−2.16	−0.19	1.97	−1.16	1.22	2.39
Belgium	4.06***	3.40***	−0.66	−0.14	1.22*	1.36
Canada	1.33*	1.54**	0.21	−0.2	0.6	0.8
Denmark	2.89**	3.18**	0.29	−0.16	2.10**	2.26
France	3.96**	3.17**	−0.79	−0.44	0.43	0.87
Germany	3.63**	3.03***	−0.59	0.62	0.97	0.35
Italy	−1.16	0.32	1.47	0.43	−1.32	−1.75
Japan	−1.55	0.44	1.99*	0.36	−0.03	−0.39
Sweden	−1.49	0.34	1.83**	−1.83	0.23	2.06
UK	1.7	1.88	0.18	1.17	1.36	0.19
USA	−5.27	−3.81	1.46**	−6.51	−0.61	5.89
Average	0.54	1.21	0.67	−0.71	0.56	1.28

This table reports the Campbell and Thompson (2008) out-of-sample R^2_{OS} (in percent), measuring the proportional reduction in mean-squared forecast error (MSFE) for the competing model relative to the historical average benchmark. All forecasts are performed on the out-of-sample period 1993-2013. The first (last) three columns present out-of-sample results without (with) imposing equity premium constraints on the coefficients. In each case, the LS column corresponds to the baseline least-square model. The baseline model is estimated with OLS (nonlinear least squares when constraints are imposed on the coefficients). The Bayes column correspond to the cross-sectional learning model. 'Δ' columns indicate the difference between baseline and Bayes forecasts so that a positive value indicates a better forecast when using Bayes coefficients. The forecasts (forecasting gains between baseline and Bayes regressions) are also evaluated according to Clark and West (2007) *MSFE-adjusted* (respectively Diebold and Mariano (1995) and West (2006)) statistic (see Section III.D). "Average" is the average of the R^2_{OS} statistics across the 11 countries.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Internet appendix

Additional results

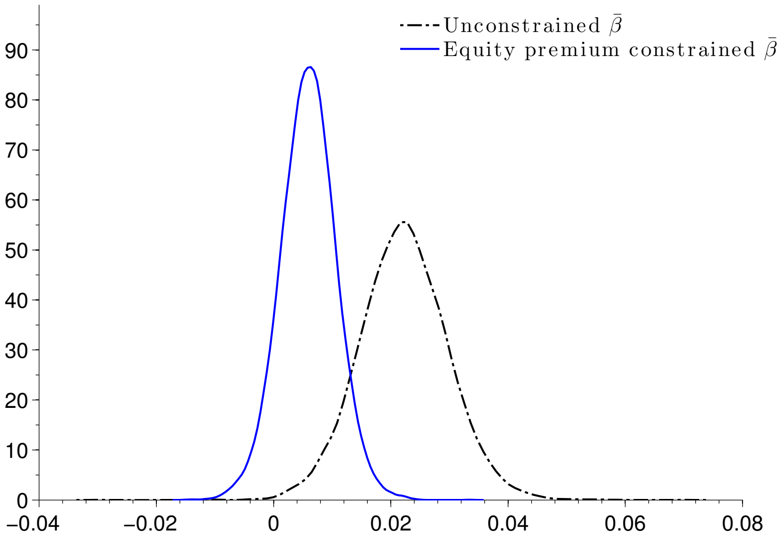

(a) Posterior distribution of $\bar{\beta}$

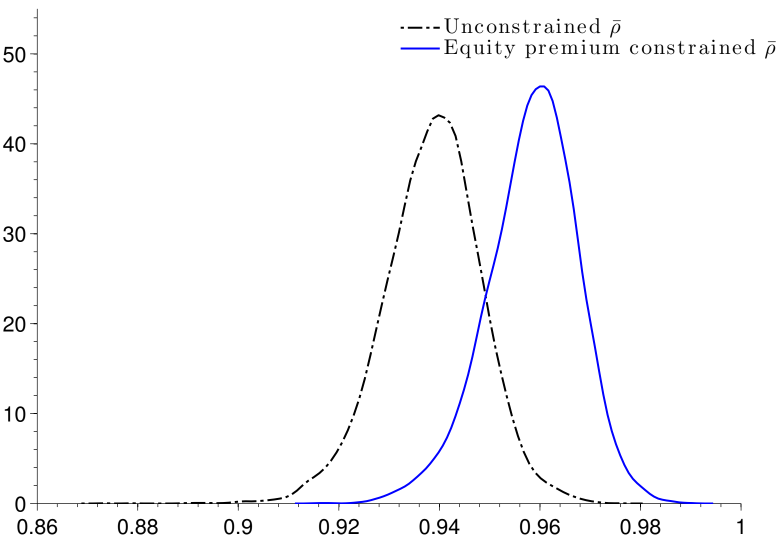

(b) Posterior distribution of $\bar{\rho}$

Figure 4.6: Log dividend-price ratio

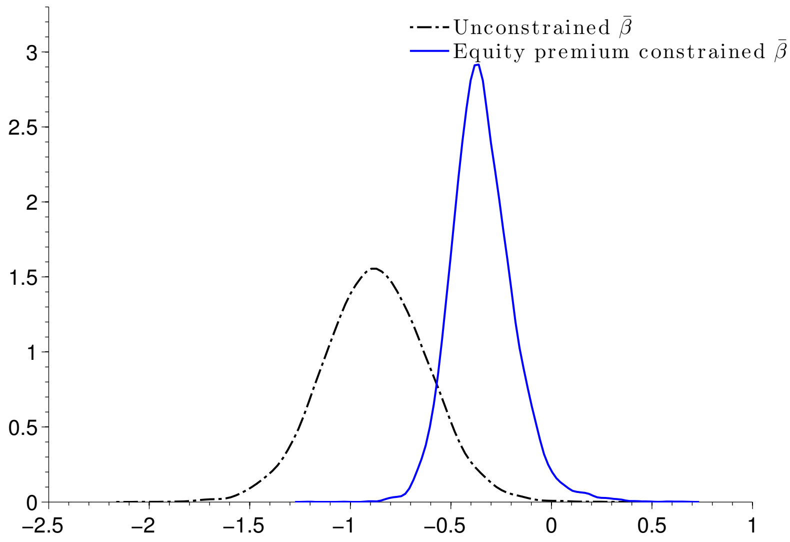

(a) Posterior distribution of $\bar{\beta}$

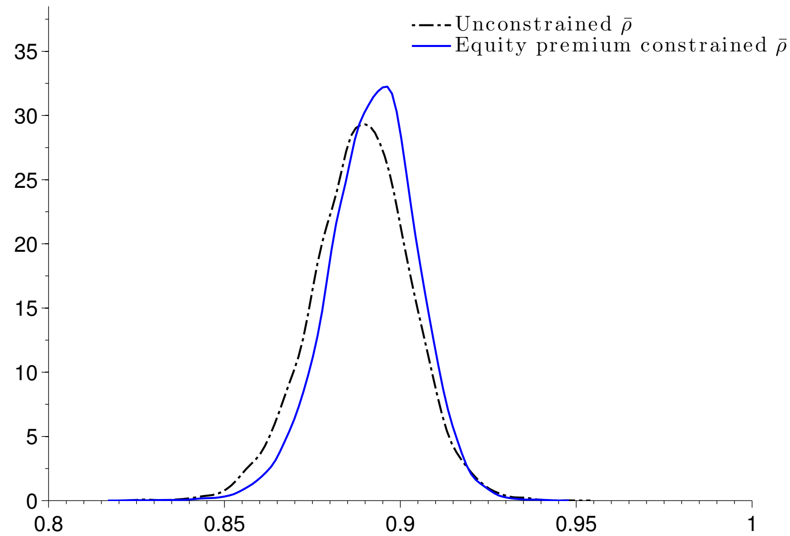

(b) Posterior distribution of $\bar{\rho}$

Figure 4.7: Short-term interest rate

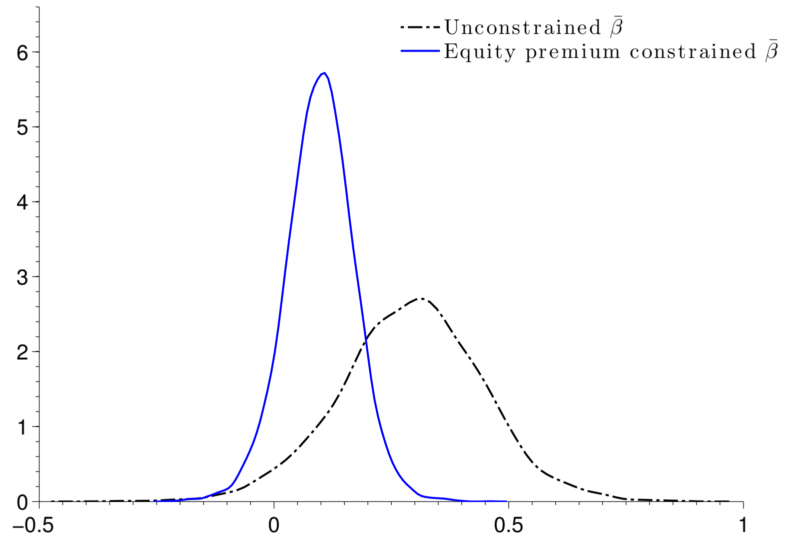

(a) Posterior distribution of $\bar{\beta}$

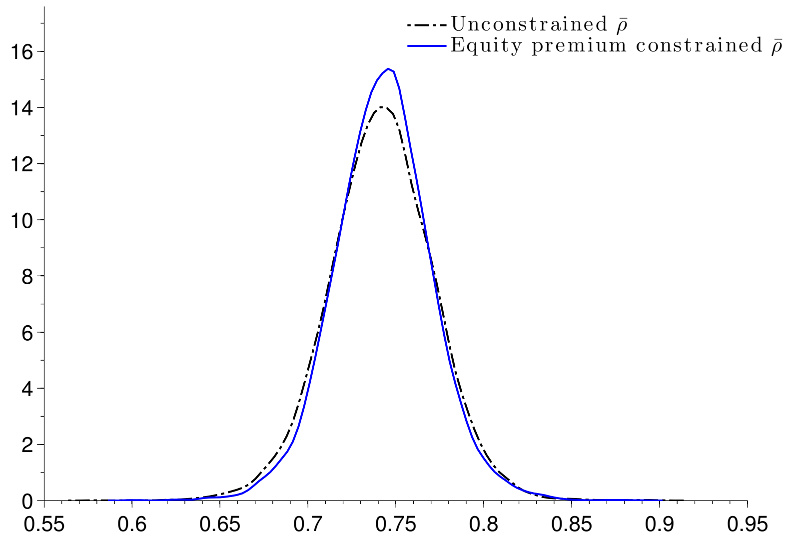

(b) Posterior distribution of $\bar{\rho}$

Figure 4.8: Term spread

Table V: Bayesian estimates: present-value model (annual frequency)

		Dividends unreinvested			Dividends reinvested		
		Excess returns		Div. growth	Excess returns		Div. growth
Prior		Neutral	EP		Neutral	EP	
Country	Sample						
Australia	Q1-1971 - Q1-2013	0.118	0.048	-0.201	0.176	0.057	-0.065
(Stdev)		(0.045)	(0.022)	(0.051)	(0.046)	(0.029)	(0.033)
$p < 0$		0.010	0.000	1.000	0.000	0.010	0.980
Belgium	Q1-1971 - Q1-2013	0.102	0.028	-0.195	0.154	0.042	-0.061
(Stdev)		(0.044)	(0.021)	(0.053)	(0.044)	(0.027)	(0.034)
$p < 0$		0.020	0.080	1.000	0.000	0.060	0.970
Canada	Q1-1971 - Q1-2013	0.089	0.026	-0.176	0.120	0.020	-0.069
(Stdev)		(0.037)	(0.015)	(0.035)	(0.038)	(0.025)	(0.032)
$p < 0$		0.010	0.050	1.000	0.000	0.200	0.990
Denmark	Q1-1972 - Q1-2013	0.091	0.025	-0.187	0.135	0.037	-0.066
(Stdev)		(0.037)	(0.015)	(0.047)	(0.035)	(0.023)	(0.031)
$p < 0$		0.010	0.050	1.000	0.000	0.060	0.990
France	Q1-1971 - Q1-2013	0.086	0.022	-0.168	0.119	0.023	-0.072
(Stdev)		(0.034)	(0.014)	(0.034)	(0.034)	(0.021)	(0.030)
$p < 0$		0.010	0.060	1.000	0.000	0.150	1.000
Germany	Q1-1971 - Q1-2013	0.106	0.024	-0.192	0.152	0.037	-0.066
(Stdev)		(0.042)	(0.019)	(0.049)	(0.042)	(0.027)	(0.034)
$p < 0$		0.010	0.100	1.000	0.000	0.100	0.980
Italy	Q1-1971 - Q1-2013	0.056	0.011	-0.230	0.122	0.031	-0.069
(Stdev)		(0.043)	(0.014)	(0.055)	(0.041)	(0.025)	(0.035)
$p < 0$		0.110	0.220	1.000	0.010	0.100	0.990
Japan	Q1-1971 - Q1-2013	0.114	0.032	-0.108	0.141	0.042	-0.047
(Stdev)		(0.036)	(0.016)	(0.030)	(0.035)	(0.021)	(0.027)
$p < 0$		0.000	0.030	1.000	0.000	0.020	0.960
Netherlands	Q1-1986 - Q1-2013	0.106	0.035	-0.190	0.163	0.053	-0.060
(Stdev)		(0.038)	(0.018)	(0.047)	(0.037)	(0.028)	(0.032)
$p < 0$		0.000	0.030	1.000	0.000	0.040	0.970
Norway	Q1-1979 - Q1-2013	0.110	0.030	-0.190	0.172	0.049	-0.060
(Stdev)		(0.044)	(0.018)	(0.054)	(0.045)	(0.032)	(0.033)
$p < 0$		0.010	0.050	1.000	0.000	0.060	0.970
Spain	Q2-1978 - Q1-2013	0.092	0.025	-0.168	0.133	0.031	-0.070
(Stdev)		(0.036)	(0.019)	(0.037)	(0.037)	(0.022)	(0.031)
$p < 0$		0.010	0.080	1.000	0.000	0.080	0.990
Sweden	Q1-1971 - Q1-2013	0.106	0.026	-0.190	0.153	0.042	-0.067
(Stdev)		(0.042)	(0.020)	(0.053)	(0.043)	(0.028)	(0.034)
$p < 0$		0.010	0.090	1.000	0.000	0.070	0.980
Switzerland	Q1-1974 - Q1-2013	0.091	0.026	-0.186	0.135	0.029	-0.066
(Stdev)		(0.041)	(0.017)	(0.055)	(0.038)	(0.026)	(0.033)
$p < 0$		0.020	0.080	1.000	0.000	0.130	0.980
UK	Q1-1971 - Q1-2013	0.179	0.070	-0.102	0.198	0.074	-0.056
(Stdev)		(0.051)	(0.024)	(0.044)	(0.045)	(0.027)	(0.032)
$p < 0$		0.000	0.000	0.990	0.000	0.000	0.960
USA	Q1-1953 - Q1-2013	0.130	0.047	-0.085	0.146	0.051	-0.053
(Stdev)		(0.038)	(0.017)	(0.037)	(0.033)	(0.023)	(0.029)
$p < 0$		0.000	0.000	0.990	0.000	0.010	0.960

This table reports Bayesian estimates for the regression $z_{i,t+1} = \mu_i + \beta_i x_{i,t} + u_{i,t+1}$ where the left hand-side variable is either excess returns or dividend growth and $x_{i,t}$ is the log dividend-price ratio. Neutral (EP) indicates that the forecasts of the equity premium are unconstrained (constrained) to be positive. The right panel corresponds to regressions where dividends are assumed to be reinvested in the stock market. Means, standard deviations, and the probability that β_i is smaller than zero are reported.

The term structure of risk

Conditional k-period Moments

This section details how to extract moments of real returns from the excess real returns in the VAR(1) model. The reader is referred to Campbell and Viceira (2004) for additional details. I drop the index i for convenience. First, I derive a set of equations relating z_{t+k} to its current value z_t plus a weighted sum of interim shocks:

$$\begin{aligned}
\mathbf{z_{t+1}} &= \Phi_0 + \Phi_1\mathbf{z_t} + \mathbf{v_{t+1}} \\
\mathbf{z_{t+2}} &= \Phi_0 + \Phi_1\mathbf{z_{t+1}} + \mathbf{v_{t+2}} \\
&= \Phi_0 + \Phi_1\Phi_0 + \Phi_1\Phi_1\mathbf{z_t} + \Phi_1\mathbf{v_{t+1}} + \mathbf{v_{t+2}} \\
\mathbf{z_{t+k}} &= \Phi_0 + \Phi_1\Phi_0 + \Phi_1^2\Phi_0 + ... + \Phi_1^{k-1}\Phi_0 + \Phi_1^k\mathbf{z_t} \\
&= +\Phi_1^{k-1}\mathbf{v_{t+1}} + \Phi_1^{k-2}\mathbf{v_{t+2}} + ... + \Phi_1\mathbf{v_{t+k-1}} + \mathbf{v_{t+k}}
\end{aligned}$$

Taking the sum and reordering terms yields:

$$\begin{aligned}
\mathbf{z_{t+1}} + ... + \mathbf{z_{t+k}} &= \left[k + (k-1)\,\Phi_1 + (k-2)\,\Phi_1^2 + ... + \Phi_1^{k-1}\right]\Phi_0 \\
&+ \left(\Phi_1^k + \Phi_1^{k-1} + ... + \Phi_1\right)\mathbf{z_t} \\
&+ \left(1 + \Phi_1 + ... + \Phi_1^{k-1}\right)\mathbf{v_{t+1}} \\
&+ \left(1 + \Phi_1 + ... + \Phi_1^{k-2}\right)\mathbf{v_{t+2}} \\
&+ ... \\
&+ (1 + \Phi_1)\,\mathbf{v_{t+k-1}} + \mathbf{v_{t+k}}
\end{aligned}$$

Or more compactly :

$$\mathbf{z_{t+1}} + \mathbf{z_{t+2}} + ... + \mathbf{z_{t+k}} = \left[\sum_{n=0}^{k-1}(k-n)\,\Phi_1^n\right]\Phi_0 + \left[\sum_{m=1}^{k}\Phi_1^m\right]\mathbf{z_t} + \sum_{q=1}^{k}\left[\sum_{p=0}^{k-q}\Phi_1^p\mathbf{v_{t+q}}\right]$$

I am now able to compute conditional variance:

$$\begin{aligned} Var\left(\mathbf{z_{t+1}}+\mathbf{z_{t+2}}+...+\mathbf{z_{t+k}}\right) &= Var\left(\left[\sum_{n=0}^{k-1}(k-n)\,\Phi_1^n\right]\Phi_0+\left[\sum_{m=1}^{k}\Phi_1^m\right]\mathbf{z_t}+\sum_{q=1}^{k}\left[\sum_{p=0}^{k-q}\Phi_1^p\mathbf{v_{t+q}}\right]\right)\\ &= Var\left(\sum_{q=1}^{k}\left[\sum_{p=0}^{k-q}\Phi_1^p\mathbf{v_{t+q}}\right]\right)\end{aligned}$$

as all other terms are constant or already known at time t. Expanding this expression yields:

$$\begin{aligned} Var\left(\mathbf{z_{t+1}}+\mathbf{z_{t+2}}+...+\mathbf{z_{t+k}}\right) &= \Sigma_i+\left(I+\Phi_1\right)\Sigma_i\left(I+\Phi_1\right)'\\ &\quad+\left(I+\Phi_1+\Phi_1\Phi_1\right)\Sigma_i\left(I+\Phi_1+\Phi_1\Phi_1\right)'\\ &\quad+...\\ &\quad+\left(I+\Phi_1+...+\Phi_1^{k-1}\right)\Sigma_i\left(I+\Phi_1+...+\Phi_1^{k-1}\right)'.\end{aligned}$$

I am only interested in extracting conditional moments per period from the portion of the VAR that contains returns, which I extract as follows:

$$E\left(r_{t\to t+k}|\Psi_i,\mathbf{z_t}\right) = \frac{1}{k}\mathbf{M}\,E\left(\mathbf{z_{t+1}}+\mathbf{z_{t+2}}+...+\mathbf{z_{t+k}}\right). \tag{1}$$

$$Var\left(r_{t\to t+k}|\Psi_i\right) = \frac{1}{k}\mathbf{M}\,Var\left(\mathbf{z_{t+1}}+\mathbf{z_{t+2}}+...+\mathbf{z_{t+k}}\right)\,\mathbf{M}'. \tag{2}$$

with $\mathbf{M}=(1,0)'$.

Predictive variance

The predictive variance is computed from the output of the MCMC. For each draw $\Psi^{(d)}$ (d=1…L) from the posterior density $p(\Psi|D)$, I compute the conditional means and variance as described above

$$E_{i,d}(k) = E\left(r_{t\to t+k}\middle|\Psi_i^{(d)},\mathbf{z_t}\right) \tag{3}$$

$$V_{i,d}(k) = Var\left(r_{t\to t+k}\middle|\Psi_i^{(d)}\right). \tag{4}$$

Observe that the conditional mean must be measured for a given $\mathbf{z_t}$. Following Hoevenaars et al. (2014), I choose to set it at the unconditional mean $\mathbf{z_t} = (I - \Phi_1)^{-1} \Phi_0$. It is then straightforward to compute the predictive variance as

$$\bar{V}_i(k) = \frac{1}{L} \sum_d V_{i,d}(k) + \frac{1}{L} \sum_d \left(E_{i,d}(k) - \bar{E}_i(k) \right)^2 \tag{5}$$

where $\bar{E}_i(k) = \frac{1}{L} \sum_d E_{i,d}(k)$.