PhD-FDEF-2014-10
The Faculty of Law, Economics and Finance

DISSERTATION

Defense held on 05/06/2014 in Luxembourg

to obtain the degree of

DOCTEUR DE L'UNIVERSITÉ DU LUXEMBOURG

EN SCIENCES FINANCIÈRES

by

Nicolas MARTELIN
Born on 26 January 1986 in Nancy, France

ESSAYS ON IRRATIONAL BEHAVIOR IN FINANCIAL MARKETS

Dissertation defense committee

Dr Thorsten Lehnert, dissertation supervisor
Professor, Université du Luxembourg

Dr Etienne Harb
Assistant Professor, ESSCA Angers

Dr Christian Wolff, Chairman
Professor, Université du Luxembourg

Dr Marie Lambert
Assistant Professor, Université de Liège

Dr Roman Kräussl, Vice Chairman
Associate Professor, Université du Luxembourg

Acknowledgements

Over the course of the four years of my doctoral education, I have been surrounded, supported and cheered up by many people to whom I would like to express my deepest gratitude for helping me complete this thesis.

First of all, I want to thank my supervisor, Prof. Thorsten Lehnert, for giving me the opportunity to embark on this journey. His valuable advice and academic expertise made the accomplishment of this work possible. I am also extremely grateful for his availability and dedication.

I would also like to thank Prof. Christian Wolff, Director of the Luxembourg School of Finance, for accepting to be on my committee and providing me with an excellent working atmosphere for conducting research. I have also benefited greatly from the insightful comments of Prof. Roman Kräussl, feedback that helped me maximize the return on my work.

My thanks go to Prof. Marie Lambert who did much to point me in right directions and increase my motivation to pursue a PhD at the University of Luxembourg. Special thanks go also to Prof. Etienne Harb who put a lot of effort into participating in my defense committee.

I am thankful to my professors from the doctoral program and particularly to Prof. Tibor Neugebauer who gave me the first insights into behavioral finance. Then there is Prof. Antonio Cosma, Prof. Jos Van Bommel, and Dr. Nikalaos Papanikolaou who all contributed constructive comments and equipped me with the financial and econometrics tools needed to address my research topic.

I want also to acknowledge the financial and technical support of the University of Luxembourg and the Luxembourg School of Finance and its staff, especially Deborah Marx and Martine Zenner who have been a precious help in making those four years go smoothly. I am also particularly grateful to Virginie Mucciante for her assistance at the beginning and at the end of my PhD and to Emmanuelle Ambroisien.

At the Luxembourg School of Finance, I also wish to acknowledge the contributions of my fellow PhD students and friends, Fanou, Lamia, Rustam, David, Mike, Magda, and Fabian, whose presence and friendship brought a softer touch to the work environment.

I give further thanks to: Laurent, who as a good friend, not only triggered my desire to pursue a doctoral education but also set high standards; Aude for not letting me be the last one to graduate; Michael for listening to my complaints; Matthias for all the fun support; Pierre for helping me think about the future; Virginie for being always there. Also, special thanks go to Yan and Sandrine, for helping me take up challenges on and off the track, as well as Gérard Gaillard, Michel Comerlatti, Dr. Thomas Schwartz Thierry Wendling, and Stephen Evans.

Finally, I would like to thank my parents, my sister, and Karin for their heartfelt encouragements and support, for sharing these years with me and keeping my spirits up along the way.

Table of Contents

Chapter One

General Introduction

As long as investors are human, they will be subject to irrationality. Prompted by emotion, their judgments get in the way of rational behavior, making decisions that drift away from what common sense would suggest. Since the seminal work of Kahneman and Tversky (1974), two schools of thought have been widely accepted as being able to describe and forecast investor behavior: the traditional rational theory, anchored in the efficient market hypothesis and the behavioral approach, which takes into account the effect of social, cognitive and emotional factors on the decision making process.

Proponents of the Efficient Market Hypothesis (rational theory) would claim that markets are fully rational and able to incorporate new information correctly into asset prices. They accept that some abnormality may arise in the formation of asset prices but assume that competition amongst investors trying to take advantage of such abnormalities will drive prices back to their "correct values". They consider Bayesian rationality as a good description of investor behavior. Therefore, we consider – in this thesis – as being "rational" any sort of behavior that is described in the theoretical background of the Efficient Market Hypothesis (EMH). In other words, we understand "rationality" as the combination of the following criteria:

1. In a given market, all participants are correct, on average (even though there can be some discrepancies amongst them, arbitrageurs will cancel them out), when it comes to pricing a financial asset and/or to incorporating new information into their future expectations.
2. Market participants have all well-defined subjective utility functions that they want to maximize and act accordingly.

However, empirical studies and the evidence of various apparently irrational speculative bubbles and crashes challenged this view and the new area of behavioral finance came into play.

Therefore, we will call "irrationality", any type of behavior that does not fulfill the criteria stated above. The term "irrationality" – here – is then to be understood as the contradiction of finance theory. It can be found in other works under the name "bounded rationality" (Simon (1997)) or "minimal rationality" (Rubinstein (2000)).

Essence of behavioral finance theory is that not all market agents are fully rational. This lack of rationality can come from either personal preference (e.g. people have a tendency to be "loss-averse" such that a loss of \$10 might make them feel bad by as much as a \$20 gain makes them feel good) or mistaken beliefs (e.g. people try to be good Bayesians but fail as cognitive biases get in the way of their rational thinking). First described by Kahneman and Tversky (1972), those biases generally lead to what is commonly called "irrationality". Ritter (2003) gathers a list of six different biases likely to be experienced by investors.

Heuristics are central to this process. They are simplified decision-making tools like rules of thumb used by investors during decision making. Though they can work much of the time, they are responsible for systematic deviations from Bayesian rationality (e.g. ignoring sampling theory).

Overconfidence is the tendency for people to misjudge their abilities or knowledge (e.g. traders experiencing the "better-than-average effect" where they believe they know more or have access to better information than the rest of the market).

Mental accounting deals with the way people frame a monetary transaction in their mind (e.g. a transaction in cash might appear of a greater value than a transaction settled with credit card).
Framing is the effect that comes from people drawing different conclusion from the same information, depending on how they are presented with this information. It leads to an inconstancy of choice.
Representativeness involves people giving excessive weight to recent news or experiences (e.g. a big rise in equity returns might appear to be "normal" if it is sustained for a certain period of time).
Conservatism refers to people being slow to pick up on changes.

Every bias listed above has been identified in several experiments and is well documented. Therefore we assume that they are stylized facts of human behavior.
Personal preferences along with cognitive biases create the behavioral heterogeneity of a given market. In this thesis we study three markets that are completely different from each other and have their own peculiarities: the stock market (US equity funds), the option market (S&P500 index options) and the art market (Impressionist and Modern, Post-war and Contemporary, American, and Latin American art indices).

In Chapter Two, we focus on how investors are prone to cognitive bias and/or personal preferences in a stock market context. We call these "noise traders." Using US equity fund flows as a proxy for noise trader sentiment, we investigate empirically if any noise trader activity affects the cross-section of S&P 500 index option prices. Inspired by previous research, we hypothesize that these daily flows in and out of mutual funds carry a substantial component of noise trader sentiment. We make use of these flows to test for its impact on a fast moving market like the options market. Other commonly used implicit or explicit measures of sentiment are typically only available at lower frequencies. Relying on those low frequency measures, previous findings suggest that market sentiment correlates contemporaneously with implied volatilities and the risk-neutral skewness, where the causality is unclear. However, volatility is obviously priced in the option market and therefore, any effect of noise trading on volatility should have an immediate affect on the cross-section of option prices. This could also be transmitted from the stock market to the options market. Daily data allow us to investigate if the cross-section of index option prices is conditional on our beginning-of-period proxy of noise trading and to test for information as well as noise transmission effects.

We find strong evidence that the noise from the stock market is transmitted to the options market. When noise traders are bearish (bullish) on a particular day, resulting in flows out of (into) US equity funds, volatilities of S&P500 index options tend to increase (decrease) significantly on the following day.

The objective of Chapter Three is to investigate a particular behavioral aspect of the options market: the overreaction puzzle identified by Stein. Stein (1989) derives and tests empirically a model that describes the relationship between implied volatilities of options of different maturities. Assuming that volatility evolves according to a continuous-time mean-reverting AR1 process, with a constant long-run mean and a constant coefficient of mean-reversion, theoretically, the implied volatility of longer maturity (two-months) options should move in a responsive, but smoothing manner to changes in implied volatility of shorter maturity (one-month) options. However, the empirical values of this elasticity exceeded the theoretical upper bound of normal-reaction. Stein interprets his findings as overreaction, which is caused by market inefficiencies, claiming that this contradicts the rational expectations hypothesis for the term structure of implied volatilities.

We show theoretically that under normal market conditions the risk-neutral volatility process is substantially more persistent than the physical one. Investors' risk aversion appears to be the main factor driving this persistence. Theoretically, long-term volatility should react more strongly to changes in short-term volatility in periods when investors are highly risk averse, and risk-neutral volatility is highly persistent. In contrast, in periods of low risk aversion, long-term volatility should react less strongly to changes in short-term volatility, because risk-neutral volatility is less persistent. Using daily data on S&P 500 index options for the 2000-2010 period, we verify empirically these theoretical predictions. In periods of high risk aversion, long-term volatility reacts strongly to changes in short-term volatility, which can be explained by highly persistent risk-neutral volatility dynamics in that period. The effect cannot be observed in periods of low risk aversion, because of less persistent volatility dynamics. Overall, we provide strong evidence that the empirical observation that Stein (1989) discovered is not overreaction, but in line with perfectly rational behavior.

In Chapter Five, we bring the study of behavioral heterogeneity to a previously unresearched setting: the art market. We answer the question raised by the record-breaking prices in the art market observed over the last three years: is there a bubble in the art market? Given the difficulty of determining the fundamental value of artwork, we apply a right-tailed unit root test with forward recursive regressions (SADF test) to detect explosive behaviors directly in the time series of four different art market segments (Impressionist and Modern, Post-war and Contemporary, American, and Latin American) from 1970 to 2013. We identify two historical speculative bubbles and find an explosive movement in the current Post-war and Contemporary and American fine art market segments.

Chapter Two

Noise Trading and the Cross-Section of Index Option Prices[i]

2.1 Introduction

This chapter investigates how option prices are affected by behavioral biases. Various researchers find evidence that arbitrage in options markets is limited and document obvious misreactions of market participants to information in the options market (see Stein (1989), Poteshman (2001), Poteshman and Serbin (2003), Mahani and Poteshman (2004), Lakonishok et al. (2006), Garleanu et al. (2007), Han (2008) and Christoffersen et al. (2010)).

Stein (1989) documents that long-term options tend to react more to changes in implied volatility than short-term options. Poteshman (2001) investigates how investors in the options market react to information in volatility changes. He finds short-horizon under-reaction and long-horizon over-reaction to daily changes in market volatility.

Poteshman and Serbin (2003) examine the early exercise of calls and find that customers of discount brokers irrationally exercise far too early. Lakonishok et al. (2006) investigate options market activity and show that during the internet bubble, investors increased purchases of calls, but not of puts and even when securities were available, option traders did not seem to have the courage to bet against the bubble.

Looking at price discovery across different markets, Berndt and Ostrovnaya (2008) find that equity option markets are more likely to trade on “unsubstantiated rumors” compared to credit derivatives markets, for example.

Bollen and Whaley (2004) examine the relation between net buying pressure and the shape of the implied volatility function for index and individual stock options. Their results suggest that changes in implied volatility are directly related to net buying pressure from public order flow. Han (2008) argues that investor sentiment is an important determinant of option prices and that observed mispricing is sentiment-driven. In line with this intuition, he finds that the index option smile is steeper (flatter) when the market is bearish (bullish). The author finds evidence that the effects appear to be more pronounced when arbitrage in index options is more limited. In other words, in times of higher of market participant uncertainty about market volatility, mistakes in attempts to forecast volatility can cause options to be mispriced. Along these lines of thinking, Frijns et al. (2010, 2012) propose an alternative, more economically motivated model to estimate volatility and price options. Heterogeneous agents form different beliefs about the future level of market volatility and they trade accordingly. The agents trade on long-term mean reversion in volatility as well as on exogenous shocks from the underlying market, but they also consider noise traders who incorporate market sentiment into their beliefs. They find that noise from the underlying market plays an important role in options markets.

The motivation for most of this type of empirical research is rooted in the theoretical work of De Long et al. (1990). The authors suggest that noise trader sentiment has an impact on the volatility of market prices. De Long et al. (1990) introduce a formal model where a positive or negative price pressure effect can exist because smart money might not be willing to trade against noise traders as the latter introduce a specific risk by their trading. They break down the overall price pressure into two parts, the “price pressure” and “hold-more” effects. According to their model, the price pressure effect is always negative, i.e. it increases prices and, thus, decreases returns. The hold-more effect describes the situation when noise traders become bullish such that they demand more of the risky asset. Thus they increase their expected returns if they hold more of the risky asset than sophisticated investors; therefore, the overall price pressure could be positive if the latter effect is dominant. On the other hand, if noise traders become bearish, the hold-more effect will also be negative, resulting in a negative overall price pressure effect. This mechanism, which prevents arbitrageurs from betting against noise traders, is the “create-space” effect: if noise trading increases volatility, price uncertainty may prevent smart investors from trading against the noise. On the other hand, the Friedman effect (also known as the “buy high - sell low” effect) just describes the situation where noise trading is exploited by smart investors, i.e. if increased volatility does not prevent smart investors from betting against noise traders. In such a case, higher volatility will translate into bigger losses for noise traders, because if they herd, their feedback trading creates even bigger losses, which will be apparent through a negative effect on returns. As before, both effects might be present at the same time, but one might dominate the other and hence determine the total effect. Unlike the broken down price pressure effect, the create-space and Friedman effect do not influence prices directly, but through an increase in volatility of assets' returns (see also e.g. Barberis et al. (1998)). Lee et al. (2002) test the implications of the model empirically and conclude that the data support the theoretical predictions. Edwards and Zhang (1998) investigate fund investors’ return-chasing behavior and price pressure effects. The authors find strong evidence for the former and some evidence for the latter[1]. They interpret this finding as support for the noise trader hypothesis and conclude that a “downward price spiral in stock prices” induced by negative fund net sales could threaten market stability and therefore affects market volatility.

The use of fund flows as a proxy for noise trader sentiment is motivated by previous findings. Data published by the Investment Company Institute shows that mutual funds are mainly retail funds and that typically retail investors (households) invest in them (ICI, 2011). Goetzmann et al. (2000) use principal component analysis of daily U.S. mutual fund flows, classified according to eight major asset classes. Their first principal component has loadings, which show a negative relationship between stock and bond

[1] See also e.g. Ippolito (1992), Sirri and Tufano (1993) and Chevalier (1997).

fund flows. The authors show that this can be interpreted as investor sentiment (see also Brown et al. (2005)). Baker and Wurgler (2006, 2007) use principal component analysis for major equity and bond classes. They show that the second principal component of flow changes has opposite loadings on speculative and safe funds flows. The authors show that this "speculative demand" is correlated to their sentiment index while the "overall demand" is not. Ben-Rephael et al. (2010) investigate a proxy for monthly shifts between bond funds and equity funds. They find that this measure is negatively correlated with changes in the VIX and is positively contemporaneously correlated with excess aggregate stock market returns. Chiu and Kini (2011) examine the relation between various phases of the equity issuance process and monthly net fund flows into equity mutual funds. Their results are also consistent with the view that equity fund flows can serve as an instrument for noise trader sentiment.

In this chapter, using US equity fund flows as a proxy for noise trader sentiment, we investigate empirically if any noise trader activity affects the cross-section of S&P 500 index option prices. Inspired by previous research, we hypothesize that these daily flows in and out of mutual funds carry a substantial component of noise trader sentiment. We make use of these flows to test for its impact on a fast moving market like the options market. Other commonly used implicit or explicit measures of sentiment are typically only available at lower frequencies. Relying on those low frequency measures, previous findings suggest that market sentiment correlates contemporaneously with implied volatilities and the risk-neutral skewness, where the causality is unclear. However, volatility is obviously priced in the option market and therefore, any effect of noise trading on volatility should immediately affect the cross-section of option prices, but could also be transmitted from the stock market to the options market. Daily data allow us to investigate if the cross-section of index option prices is conditional on our beginning-of-period proxy of noise trading and to test for information as well as noise transmission effects. The remainder of this chapter is organized as follows: in Section 2.2, we discuss the data, in Section 2.3, we present the empirical analysis and results and section 2.4 concludes.

2.2 Data

Fund Flows

We use mutual fund flows of domestic US equity funds as a proxy for investor sentiment (see Ben-Rephael et al. (2010) and Chiu and Kini (2011)). Daily aggregate mutual fund flows are provided by TrimTabs for the period January 1998 to December 2004. The data set is based on information from, on average, 1625 funds that report net asset values and total net assets to TrimTabs on a daily basis. We apply a series of rigorous checks and filters to remove any kind of error in the data (see Chalmers et al. (2001), Greene and Hodges (2002) and for a detailed discussion see Beaumont et al. (2008)).

From these numbers, daily flows in and out of mutual funds are calculated for each mutual fund. Finally, we only retain flows of domestic equity funds (706 funds) according to Wiesenberg and Morningstar information. Based on these flows, we construct equally and value-weighted average flows per day[2]. Results suggest that both series are stationary, but show some non-normality and autocorrelation for up to three lags. Interestingly, the equally-weighted flows exhibit only minor negative autocorrelation and the lowest non-normality. Subsequently, we use the equally-weighted flows as our proxy for noise trader activity in the analysis. However, results turn out to be robust regardless of the choice of the proxy for sentiment.

S&P500 Options

We follow common practice in the empirical option pricing literature and use European options on the S&P 500 index (symbol: SPX). The market for S&P index options and futures is the most active index options and futures market in the world. The last trading day is the third Friday of the expiration month if that is an exchange trading day; otherwise, it is the first possible day prior to that Friday.

[2] Summary statistics of the flow data are reported in Beaumont et al. (2008). For better comparison, we scale the flows by a factor 1/100.

We follow Barone-Adesi et al. (2008) in filtering the original options data. For liquidity reasons, we only consider closing prices of out-of-the money put and call SPX options for each trading day. We obtain options data from OptionMetrics for the period January 2000 – December 2004. The bid-ask midpoint prices are taken. In line with Bollen and Whaley (2004), we exclude options with absolute call deltas below 0.02 or above 0.98 because of distortions caused by price discreteness. The delta of a European-style call option is

$$\text{delta}_C = N\left(\frac{\ln((S - PVD)e^{rT} / K) + 0.5\sigma^2 T}{\sigma\sqrt{T}}\right), \tag{2.1}$$

where S-PVD is the dividend-adjusted index level, K is the exercise, σ is the implied volatility of the option, r is the risk-free interest rate corresponding to the time to maturity (T) of the option. Table 2.1 presents the moneyness categories according to Bollen and Whaley (2004) that we use for the subsequent analysis.

Table 2.1 Definition of Moneyness Categories (Bollen and Whaley (2004)).
This table shows five different moneyness categories defined as in Bollen and Whaley (2004) on the delta of each option. We exclude options with absolute deltas below 0.02 or above 0.98. Daily option data are from OptionMetrics and cover the period January 2000 to December 2004.

Category	Labels	Range
1	Deep in-the-money (DITM) call Deep out-of-the-money (DOTM) put	$0.875 < \Delta c \leq 0.98$ $-0.125 < \Delta p \leq -0.02$
2	In-the-money (ITM) call Out-of-the-money (OTM) put	$0.625 < \Delta c \leq 0.875$ $-0.375 < \Delta p \leq -0.125$
3	At-the-money (ATM) call At-the-money (ATM) put	$0.375 < \Delta c \leq 0.625$ $-0.625 < \Delta p \leq -0.375$
4	Out-of-the-money (OTM) call In-the-money (ITM) put	$0.125 < \Delta c \leq 0.375$ $-0.875 < \Delta p \leq -0.625$
5	Deep out-of-the-money (DOTM) call Deep in-the-money (DITM) put	$0.02 < \Delta c \leq 0.125$ $-0.98 < \Delta p \leq -0.875$

The underlying S&P 500 index level, dividend yields and the term structure of zero-coupon default-free interest rates are also provided by OptionMetrics. On each day, we fit a functional form with curvature to the term structure in order to obtain the interest rate that matches the maturity of the option. We price the options by using the dividend-adjusted underlying S&P level. We calculate the daily Put-Call ratio by dividing the volume of all out-of-the money put options by the volume of all out-of-the money call options that fulfill our criteria on a particular day.

Summary statistics

Table 2.2 summarizes the trading activity in S&P 500 index options over the January 2000 to December 2004 sample period. The total number of contracts traded in each moneyness category is reported. Overall, nearly 86 million out-of-the-money contracts were traded in the 5 year-period, roughly one-third of them are calls and two-thirds are puts. Comparing across moneyness categories, trading volume for calls is greatest for out-of-the money calls (category 4), followed by deep out-of-the money calls (category 5). Relatively symmetric, trading volume for puts is greatest for out-of-the money puts (category 2), followed by deep out-of-the money puts (category 1). Around 20% of the trading volume is in at- or near-the-money options, either slightly out-of-the money calls or slightly out-of-the money puts. This evidence is consistent with the use of S&P 500 index puts as portfolio insurance by equity portfolio managers. Given that we look at the post internet bubble period, it is not surprising that out-of-the money put options dominate our sample. Table 2.2 contains the average implied volatilities of the S&P 500 index options over the period January 2000 to December 2004. As the results in the table show, the index implied volatilities decrease monotonically across the delta categories.

Table 2.2 Summary of Number of Contracts and of Implied Volatilities.

Daily option data are from OptionMetrics and cover the period January 2000 to December 2004. The first column shows the corresponding delta moneyness category defined in Table 1. Column two and four show the number of SPX contracts traded during that period for calls and puts, respectively. Columns three and five show the corresponding proportions. We exclude deep in-the money (DITM) and in-the-money (ITM) calls and puts from our sample (denoted by --). The bid-ask midpoint of option prices and the dividend-adjust underlying S&P 500 are taken to estimate implied volatilities using the analytical Black-Scholes-Merton formula. The last two columns at the right show the average implied volatility of each moneyness category and the average difference between implied and realized volatility. As in Bollen and Whaley (2004), the realized volatility of the last 60 trading days prior to each measurement of implied volatility is computed.

Delta Value Category	Calls		Puts		Average Implied Volatility	Average Diff. between Implied and Realized Volatility
	No. of Contracts	Prop. of Total	No. of Contracts	Prop. of Total		
1	--	--	19,355,307	0.226	0.284	0.092
2	--	--	25,982,293	0.303	0.229	0.036
3	6,694,369	0.078	10,980,298	0.128	0.196	0.003
4	15,170,413	0.177	--	--	0.174	-0.018
5	7,636,768	0.089	--	--	0.166	-0.026
Totals	29,501,550	0.344	56,317,898	0.656		

The average implied volatility of the category 1 options (deep out-of-the-money puts) is 28.4 percent, around 12 percentage points higher than the average implied volatility of category 5 options (deep out-of-the-money calls), 16.6 percent. Figure 1 presents the time series of S&P500 returns and flows (scaled).

Figure 2.1 Daily S&P 500 returns and scaled flows.

This figure shows daily S&P 500 returns and scaled flows. Daily aggregate mutual fund flows are from TrimTabs and cover the period January 2000 to December 2004. A series of rigorous checks and filters are applied to remove any kind of error in the data. The sample consists of all flows of domestic equity funds (706 funds) according to Wiesenberg and Morningstar information. Scaled value weighted flow is the aggregated, value-weighted average flow per day where each individual flow has been normalized by the size of the corresponding fund.

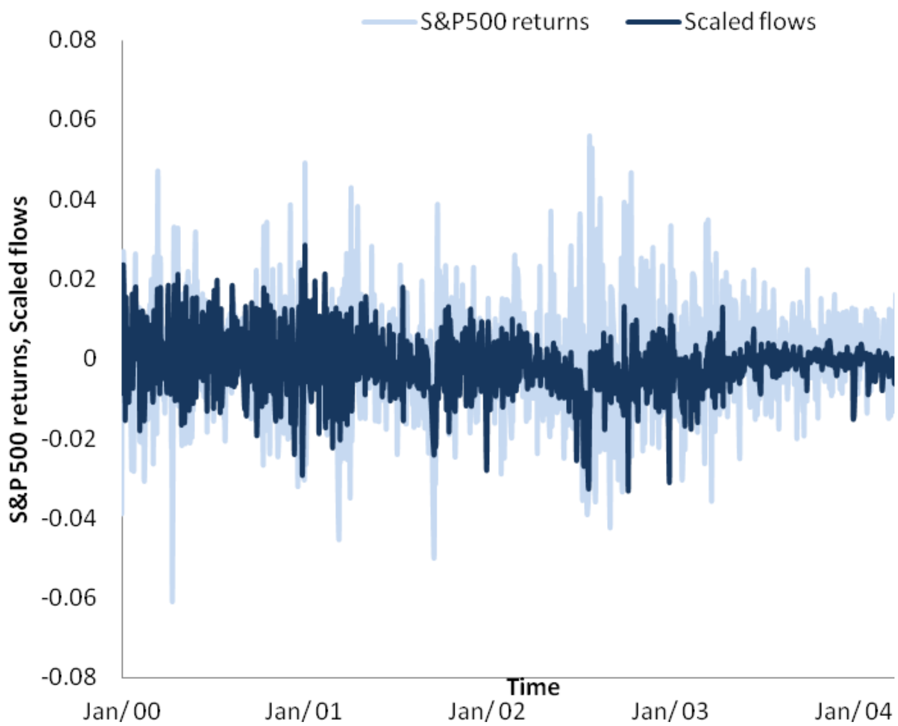

2.3 Empirical Analysis

Estimating the implied volatility surface

For the empirical analysis, first we use a modification of the prominent ad-hoc Black-Scholes model of Dumas, Fleming and Whaley (1998) to estimate the implied volatility surface of index options. The aim is to use all available information content in index option prices and to investigate the time series of standardized implied volatilities for fixed moneyness options with fixed time to maturities. Rather than averaging the two contracts that are closest to at-the-money or closest to one-month maturity, we fit a modified ad-hoc Black-Scholes model to all option contracts on a given day and subsequently obtain the desired implied volatility. This strategy eliminates successfully some of the noise from the data (see

Christoffersen et al. (2010)). As indicated in Bollen and Whaley (2004), it is industry practice to quote Black-Scholes volatilities by option delta. Therefore, we allow each option to have its own Black-Scholes implied volatility depending on the call options delta and time to maturity T. We use the following functional form for the options implied volatility[3]:

$$IV_{i,j} = \alpha_0 + \alpha_1 \text{delta}_{C i,j} + \alpha_2 \text{delta}_{C i,j}^{2} + \alpha_3 T_j + \alpha_4 T_j^2 + \alpha_5 \text{delta}_{C i,j} T_j, \tag{2.2}$$

where IV_{ij} denotes the implied volatility and $\text{delta}_{C,i,j}$, the delta of a call option for the i-th exercise price and j-th maturity, defined in Equation (2.1)[4]. T_j denotes the time to maturity of an option for the j-th maturity. It is common practice to estimate the parameters using standard OLS. For every call option delta and maturity, we can compute the implied volatility and derive option prices using the Black-Scholes model.

For example, the implied volatility for an ATM short-term call option with one-month maturity can be derived by setting delta equal to 0.5 and time to maturity T equal to 1/12. We classify all options into five categories according to the criteria described in Bollen and Whaley (2004) (see Table 2.1). When presenting our findings, we focus on at-the-money options and deep out-of-the-money calls and puts, Bollen and Whaley (2004) moneyness categories 3, 5 and 1, respectively. Figure 2.2 compares daily changes in at-the-money (delta_C=0.5) short term (1 month) implied volatilities (category 3) with fund flows over the period January 2000 through December 2004.

[3] We have also tried other specifications for the functional form that are frequently used in the literature (replacing delta by exercise price K or moneyness K/F, where F is the forward rate, but the results are robust to these changes.
[4] For put options, we use the corresponding call delta in the implied volatility regression.

Figure 2.2 Daily changes in ATM implied volatility and scaled flows.

This figure shows daily changes in short term (1 month) ATM Implied Volatility and normalized value-weighted flows. Daily option data are from OptionMetrics and cover the period January 2000 to December 2004. The bid-ask midpoint of option prices and the dividend-adjust underlying S&P 500 are taken to estimate implied volatilities using the analytical Black-Scholes-Merton formula. ATM options correspond to the third moneyness category defined in Table 1. Daily aggregate mutual fund flows are from TrimTabs and cover the period January 2000 to December 2004. A series of rigorous checks and filters are applied to remove any kind of error in the data. The sample consists of all flows of domestic equity funds (706 funds) according to Wiesenberg and Morningstar information. Scaled value weighted flow is the aggregated, value-weighted average flow per day where each individual flow has been normalized by the size of the corresponding fund.

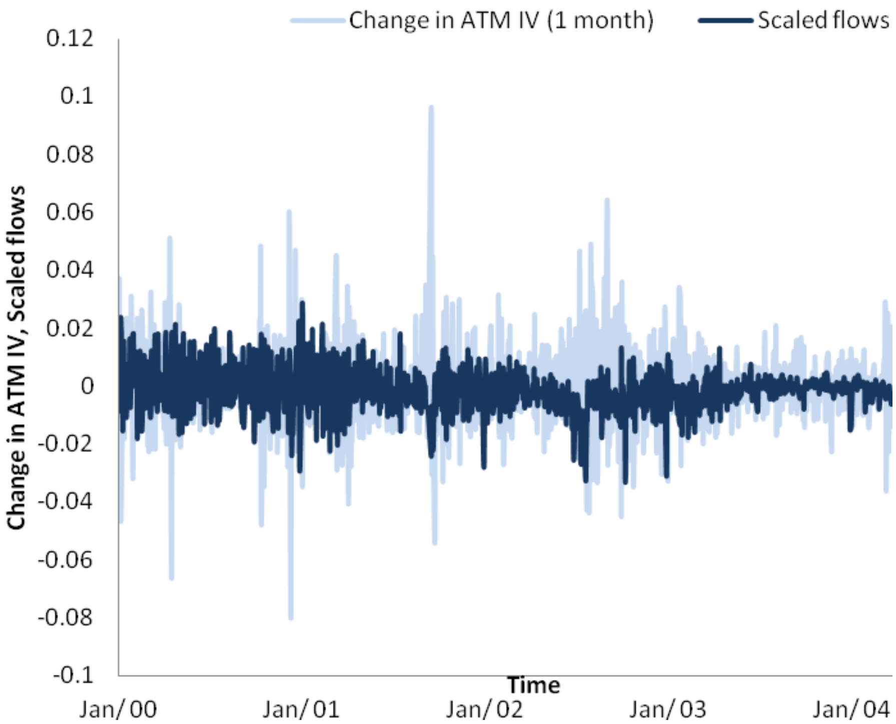

Calculating risk-neutral moments

In order to characterize the dynamics in the option market, we break down the cross-section of index option prices into the moments of the risk-neutral distribution. Bakshi et al. (2003) derive model-free measures of risk-neutral variance ($Var_t(T)$), skewness ($Skew_t(T)$) and kurtosis ($Kurt_t(T)$) based on all options over the complete moneyness range for a particular time to maturity T,

$$\mathrm{Var}_t(T) = e^{rT} V_t(T) - \mu_t^2(T)$$

$$\mathrm{Skew}_t(T) = \frac{e^{rT} W_t(T) - 3\mu_t(T) e^{rT} V_t(T) + 2\mu_t(T)^3}{\left[e^{rT} V_t(T) - \mu_t(T)^2\right]^{\frac{3}{2}}}$$

$$\mathrm{Kurt}_t(T) = \frac{e^{rT} X_t(T) - 4\mu_t(T) e^{rT} W_t(T) + 6 e^{rT} \mu_t(T)^2 V_t(T) - 3\mu_t(T)^4}{\left[e^{rT} V_t(T) - \mu_t(T)^2\right]^2} \tag{2.3}$$

where

$$\mu_t(T) = e^{rT} - 1 - \frac{e^{rT}}{2} V_t(T) - \frac{e^{rT}}{6} W_t(T) - \frac{e^{rT}}{24} X_t(T)$$

$$\mathrm{V}_t(T) = \int_{S_t}^{\infty} \frac{2(1 - \ln(\frac{K}{S_t}))}{K^2} C_t(T,K) dK + \int_0^{S_t} \frac{2(1 + \ln(\frac{S_t}{K}))}{K^2} P_t(T,K) dK$$

$$\mathrm{W}_t(T) = \int_{S_t}^{\infty} \frac{6\ln(\frac{K}{S_t})) - 3\left[\ln(\frac{K}{S_t})\right]^2}{K^2} C_t(T,K) dK - \int_0^{S_t} \frac{6\ln(\frac{S_t}{K})) - 3\left[\ln(\frac{S_t}{K})\right]^2}{K^2} P_t(T,K) dK$$

$$\mathrm{X}_t(T) = \int_{S_t}^{\infty} \frac{12\ln(\frac{K}{S_t})) - 4\left[\ln(\frac{K}{S_t})\right]^3}{K^2} C_t(T,K) dK + \int_0^{S_t} \frac{12\ln(\frac{S_t}{K})) - 4\left[\ln(\frac{S_t}{K})\right]^3}{K^2} P_t(T,K) dK$$

The parameters correspond to the ones used in Equation (2.1). C and P refer to call and put prices. Again, rather than averaging the implied volatilities of all contracts that are closest to one particular maturity (e.g. 1 month), we derive the Bakshi et al. (2003) moments using the estimated implied volatility surface and the corresponding call and put prices.

This strategy successfully ensures that the time series of risk-neutral moments is less noisy and contains fewer outliers compared to, e.g., the CBOE S&P 500 Skew index or the monthly estimates presented in Han (2008). We also use different, more ad-hoc measures of skewness and kurtosis derived from the implied volatility surface of index option prices. It is industry practice to use a risk reversal measure as a proxy for skewness and a butterfly spread measure as a proxy for kurtosis of the implied volatility surface of index option prices. Risk reversals are the difference between implied volatilities of deep out-of-the-money calls and deep out-of-the-money puts. The more negative the measure, the more pronounced is the index option smirk. Butterfly spreads are the average of implied volatilities of deep out-of-the-money calls and deep out-of-the-money puts minus the implied volatility of at-the-money options. Overall, given that the ad-hoc measures are only proxies, we expect the Bakshi et al. (2003) model-free risk-neutral moments to be more reliable compared to the ad-hoc measures.

Figure 2.3 shows the daily S&P 500 index (scaled), the risk-neutral skewness of short term (one month) options calculated according to Bakshi et. al. (2003) and the implied volatility of ATM short term (one month) options. As generally found, the skewness for S&P 500 index options is always negative through our sample period, with its minimum on December 23rd 2003 and its peak in March 2003. Figure 2.4 presents the risk-neutral skewness vis-à-vis the risk reversal measure and the risk-neutral kurtosis vis-à-vis the butterfly spread measure. The skewness and kurtosis are computed as in Bakshi et al. (2003) and the risk reversal measure as well as the butterfly-spread measure are calculated from the daily implied volatility surface of index options. While the ad-hoc measures are based on only point estimates of the implied volatility surface, the risk-neutral moments are based on all put and call prices over the whole delta range. It can be seen that the dynamics of the proxies (risk reversal butterfly spread) are quite different from the dynamics of the actual measures (risk-neutral skewness and kurtosis), which indicates that only relying on the rough proxies would be inappropriate for our subsequent time-series analysis.

Figure 2.3 S&P 500, Bakshi et al. (2003) Skewness and ATM IV.

This figure shows the S&P 500 (scaled by its beginning of the sample period value in January 2000, risk-neutral skewness for short-term (1 month) options and the implied volatility of short term (1 month) ATM S&P 500 options from January 2000 to December 2004. The skewness is computed as in Bakshi et al. (2003), and ATM options are those of category 3 defined in Table 1.

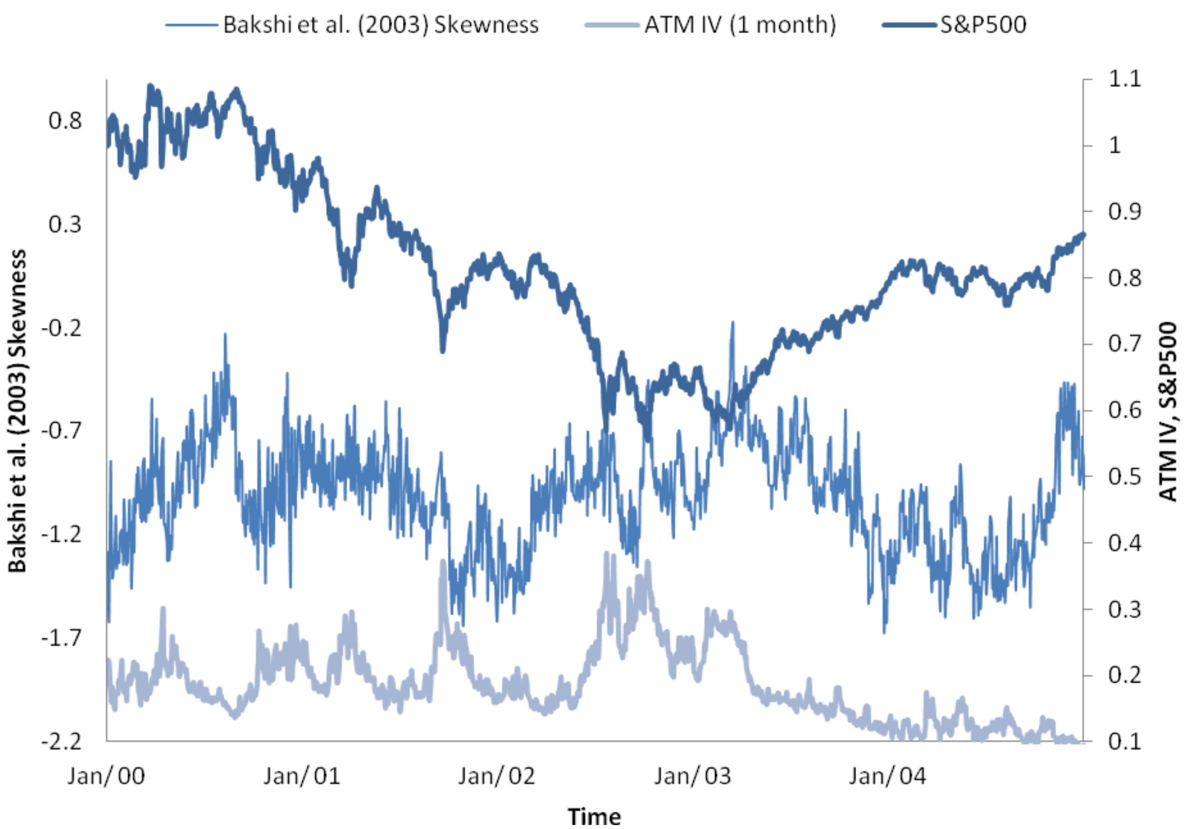

Extracting news and noise

Previous research suggests that there is a contemporaneous relationship between sentiment proxies and index returns, but the causality is not well understood. Since aggregated US equity fund flows and S&P 500 index returns on a daily frequency can be assumed to be highly dependent, this would necessarily be a source of distortion for our subsequent analysis. In our view, Bollen and Whaley (2004) and Han (2008) fail to address this issue. Additionally, it is well understood that option markets and stock markets are highly correlated. Therefore, in order to test for information transmission from the stock market into the options market, we have to control for noise that is transmitted from fund flows into the stock market as well as information that is transmitted from the option market into the stock market. Therefore, the first step in our analysis is to extract the noise component from our flow data. We regress daily fund flows on contemporaneous and lagged returns, contemporaneous and lagged changes in volatility[5] as well as on lagged fund flows in order to extract the residual component,

$$\mathrm{F}lows_t = \varpi_0 + \sum_{k=1}^{5} \psi_k \mathrm{F}lows_{t-k} + \sum_{k=0}^{5} \upsilon_k \mathrm{R}eturns_{t-k} + \sum_{k=0}^{5} \omega_k \Delta Vol_{t-k} + \mathrm{F}I_t \tag{2.4}$$

We do this for up to five lags to absorb any contemporaneous information transmission and any lagged information or noise transmission. In this way, we are able to identify the noise component of fund flows, which is not based on information that has been revealed in the stock or option market. The resulting residuals $\mathrm{F}I_t$ can be interpreted as noise signals from fund flows. This approach allows us to subsequently isolate the transmission effect of noise trading on option prices that is not a result of the price effects of noise trading. To also control for higher order autocorrelation in the returns and options market related effects, we also regress daily returns on contemporaneous and lagged changes in volatility and lagged returns for up to five lags,

$$\mathrm{R}eturns_t = \varpi_0 + \sum_{k=1}^{5} \lambda_k \mathrm{R}eturns_{t-k} + \sum_{k=0}^{5} \psi_k \Delta Vol_{t-k} + \mathrm{R}I_t \tag{2.5}$$

To also control for higher order autocorrelation in the volatility, we also regress daily changes in volatility on lagged changes in volatility for up to five lags,

$$\Delta Vol_t = \varpi_0 + \sum_{k=1}^{5} \psi_k \Delta Vol_{t-k} + \mathrm{V}I_t \tag{2.6}$$

[5] We use at-the-money implied volatility or the volatility of the risk-neutral distribution depending on the subsequent regression analysis (ad-hoc or based on moments).

Only the residuals (e.g. RI, FI and VI) from the previous regressions are used for the subsequent analysis, which allows us to isolate the two direct effects of information and noise transmission into the options market. A motivation and detailed discussion of the usefulness of this approach for testing transmission effects can be found in Acharya and Johnson (2007) and Berndt and Ostrovnaya (2008).

Regression results

In the following section, we analyze the impact of noise and information on the cross-section of index option prices. Firstly, we use an ad-hoc approach, where the key variables are derived from the observed implied volatility surface. The variables are the at-the-money implied volatility, the out-of-the-money put and call implied volatilities, a risk reversal measure used as a proxy for skewness and a butterfly spread measure used as a proxy for kurtosis. Secondly, we break down the cross-section of index option prices into the moments of the risk-neutral distribution. The moments are the risk-neutral variance, skewness and kurtosis.

Ad-hoc analysis using implied volatilities

In particular, using similar time-series model specification as Bollen and Whaley (2004), we test for lagged relationships and run the following regressions:

$$\Delta IV_t^{ATM,T} = \beta_0 + \beta_1 \Delta IV_{t-1}^{ATM,T} + \beta_2 \mathrm{F}I_{t-1} + \beta_3 \mathrm{R}I_{t-1} + \beta_4 \Delta \frac{PVo}{CVo}_{t-1} + \varepsilon_t \tag{2.7}$$

$$\Delta IV_t^{OTMp,T} = \gamma_0 + \gamma_1 \Delta IV_{t-1}^{OTMp,T} + \gamma_2 \mathrm{F}I_{t-1} + \gamma_3 \mathrm{R}I_{t-1} + \gamma_4 \Delta \frac{PVo}{CVo}_{t-1} + \varepsilon_t \tag{2.8}$$

$$\Delta IV_t^{OTMc,T} = \delta_0 + \delta_1 \Delta IV_{t-1}^{OTMc,T} + \delta_2 \mathrm{F}I_{t-1} + \delta_3 \mathrm{R}I_{t-1} + \delta_4 \Delta \frac{PVo}{CVo}_{t-1} + \varepsilon_t \tag{2.9}$$

$$\Delta RR_t^{T} = \delta_0 + \delta_1 \Delta RR_{t-1}^{T} + \delta_2 \mathrm{F}I_{t-1} + \delta_3 \mathrm{R}I_{t-1} + \delta_4 \mathrm{V}I_{t-1} + \delta_5 \Delta \frac{PVo}{CVo}_{t-1} + \varepsilon_t \tag{2.10}$$

$$\Delta BFS_t^T = \delta_0 + \delta_1 \Delta BFS_{t-1}^T + \delta_2 \mathrm{F}I_{t-1} + \delta_3 \mathrm{R}I_{t-1} + \delta_4 \mathrm{V}I_{t-1} + \delta_5 \Delta \frac{PVo}{CVo}_{t-1} + \varepsilon_t \quad (2.11)$$

where $\Delta IV_t^{ATM,T}$ are daily changes in implied volatilities of at-the-money options with time to maturity T (category 3, delta$_c$=0.5), $\Delta IV_t^{OTMp,T}$ are daily changes in implied volatilities of deep out-of-the-money put options with time to maturity T (delta$_c$=0.9275, the average of category 1) and $\Delta IV_t^{OTMc,T}$ are daily changes in implied volatilities of deep out-of-the-money call options with time to maturity T (delta$_c$=0.0725, the average of category 5) (see Table 1). ΔRR_t^T are daily changes in the risk reversal measure defined as $RR_t = IV_t^{OTMc,T} - IV_t^{OTMp,T}$ and ΔBS_t^T are daily changes in the butterfly spread measure defined as $BFS_t = (IV_t^{OTMc,T} + IV_t^{OTMp,T})/2 - IV_t^{ATM,T}$. Implied volatilities have been calculated according to Equation (2). FI_t, RI_t and VI_t are the innovations in fund flows, returns and changes in volatility, respectively and $\Delta \frac{PVo}{CVo}_t$ are daily changes in the ratio of all traded put contracts and call contracts.

In Table 2.3, we present the regression results for short (Panel A) and long term (Panel B) options (T=1 and 6 months). The dependent variables are the changes in implied volatility of at-the-money options, deep out-of-the-money puts and deep out-of-the-money calls, respectively, as well as changes in the risk reversal and butterfly spread measure. Independent variables are always the lagged dependent variable and other explanatory variables. All explanatory variables enter the regression equation as lagged variables.

Table 2.3 : Regression Results (Ad-hoc Measures)

Panel A: 1 months

	ATM IVs		OTM put IVs		OTM call IVs	
	Coefficient	T-stat	Coefficient	T-stat	Coefficient	T-stat
Constant	-0.00	-0.16	-0.00	-0.19	-0.00	-0.15
Dependent $_{-1}$	-0.09***	-3.27	-0.10***	-3.52	-0.19***	-6.54
Returns $_{-1}$	-0.08*	-1.79	-0.12*	1.91	0.10**	2.15
Flows $_{-1}$	-0.21***	-4.17	-0.28***	-4.21	-0.13***	-2.84
Put/Call ratio $_{-1}$	0.00	0.80	0.00	0.14	0.00	1.15
R^2	2.5%		2.7%		5.1%	

	Risk-Reversals		Butterfly Spread	
	Coefficient	T-stat	Coefficient	T-stat
Constant	0.00	0.13	0.00	-0.08
Dependent $_{-1}$	-0.41***	-13.73	-0.35***	-12.81
Returns $_{-1}$	0.11*	1.96	0.06***	3.27
Flows $_{-1}$	0.14**	2.49	0.01	0.43
Volatility $_{-1}$	-0.12***	-3.40	0.03**	2.56
Put/Call ratio $_{-1}$	0.00	0.74	0.00	0.16
R^2	15.8%		16.0%	

Panel B: 6 months

	ATM IVs		OTM put IVs		OTM call IVs	
	Coefficient	T-stat	Coefficient	Coefficient	T-stat	Coefficient
Constant	-0.00	-0.33	-0.00	-0.39	-0.00	-0.23
Dependent $_{-1}$	-0.05	-1.60	-0.15***	-5.41	-0.35***	-13.30
Returns $_{-1}$	-0.05***	-2.77	-0.06**	-1.99	-0.01	-0.40
Flows $_{-1}$	-0.05**	-2.07	-0.04	-1.10	-0.04	-1.20
Put/Call ratio $_{-1}$	0.00*	1.79	0.00	1.51	0.00	0.54
R^2	1.4%		2.7%		12.7%	

	Risk-Reversals		Butterfly Spread	
	Coefficient	T-stat	Coefficient	T-stat
Constant	0.00	0.24	-0.00	-0.07
Dependent $_{-1}$	-0.47***	-17.91	-0.39***	-14.01
Returns $_{-1}$	0.09***	2.89	0.00	0.05
Flows $_{-1}$	-0.01	-0.25	0.02	0.91
Volatility $_{-1}$	-0.22***	-4.59	0.00	0.14
Put/Call ratio $_{-1}$	-0.00	-1.08	0.00	0.37
R^2	20.7%		14.9%	

Notes: This table shows the regression results of time series regressions according to equations (7-11), where ad-hoc measures of the implied volatility surface of S&P500 index options are regressed on lagged explanatory variables. In panel A, we present the results for short term options (T=1 month) and in panel B for long term options (T=6 months). The first row shows the dependent variable of each pair of regressions and the explanatory variables are listed in the first column. We report the respective estimate and t-value, where standard errors are adjusted for heteroscedasticity and serial correlation according to Newey and West (1987). *** indicates statistical significance at the 1 percent level, ** at the 5 percent level and * at the 10 percent level.

Looking at the results for short-term options (Panel A), it is apparent from the Table that there is typically a significant negative serial correlation in changes in implied volatility, also documented in other studies. Regression results further suggest that there is typically a weak significant relationship between previous day news from the stock market (proxied by the residuals from the return regressions) and changes in volatility. This is evidence that not all information in the stock market is immediately incorporated into option prices, resulting in a lagged relationship between returns and implied volatilities. Therefore, we find late reaction of the option market towards news from the stock market. However, on the other hand, the lagged relationship between noise (proxied by the residuals from the flow regressions) and changes in volatility is substantial and strongly significant. In other words, if mutual fund / retail investors become bearish (bullish) on one day, implied volatilities of index options tend to increase (decrease) on the following day. The impact on implied volatilities is stronger for out-of-the-money puts compared to at-the-money options and out-of-the-money calls, and for short term options compared to long term options. Values for R squared are in line with the results presented in Bollen and Whaley (2004). The results of the other control variable, the ratio of all traded put contracts relative to all traded call contracts, are typically insignificant.

The results related to the two other measures, the risk-reversal and butterfly spread measure, are also presented in Table 2.3. Changes in stock market volatility (proxied by the residuals from the volatility regressions in Equation (2.6) affect the risk reversals negatively, suggesting that an increase in market volatility also increases the spread between out-of-the-money call and out-of-the-money put volatility and, therefore, leads to a more pronounced implied volatility smirk. Again, all information from the stock market does not seem to be incorporated already in option prices, leading to a significant relationship of lagged return (news) innovations and risk reversals. However, noise also seems to affect option prices, resulting in a significant positive relationship of flow innovations and risk reversals. In other words, if mutual fund / retail investors become bearish (bullish) on one day, the implied volatility smirk becomes more (less) pronounced on the following day. In contrast, only information from the stock market seems to affect the spread between out-of-the-money and at-the-money implied volatilities, indicating that positive (negative) returns lead to an increase (decrease) of the butterfly spread. The results are typically more pronounced for short term options compared long-term options.

Analysis using risk-neutral moments

Additionally, we test for the impact of noise and information on the risk-neutral moments derived from the risk-neutral distribution of option prices discussed in an earlier section of the paper. In particular, the regression model specification is

$$\Delta Vol_t^T = \theta_0 + \theta_1 \Delta Vol_{t-1}^T + \theta_2 \mathrm{F}I_{t-1} + \theta_3 \mathrm{R}I_{t-1} + \theta_5 \Delta \frac{PVo}{CVo}_{t-1} + \varepsilon_t, \tag{2.12}$$

$$\Delta Skew_t^T = \theta_0 + \theta_1 \Delta Skew_{t-1}^T + \theta_2 \mathrm{F}I_{t-1} + \theta_3 \mathrm{R}I_{t-1} + \theta_4 \mathrm{V}I_{t-1} + \theta_5 \Delta \frac{PVo}{CVo}_{t-1} + \varepsilon_t, \tag{2.13}$$

$$\Delta Kurt_t^T = \theta_0 + \theta_1 \Delta Kurt_{t-1}^T + \theta_2 \mathrm{F}I_{t-1} + \theta_3 \mathrm{R}I_{t-1} + \theta_4 \mathrm{V}I_{t-1} + \theta_5 \Delta \frac{PVo}{CVo}_{t-1} + \varepsilon_t, \tag{2.14}$$

where ΔVol_t^T, $\Delta Skew_t^T$ and $\Delta Kurt_t^T$ are the changes in risk-neutral volatility, skewness and kurtosis for options with time to maturity T, as defined earlier. $\mathrm{F}I_t$, $\mathrm{R}I_t$ and $\mathrm{V}I_t$ are, respectively, the innovations in fund flows, returns and volatility. As well, $\Delta \frac{PVo}{CVo}_t$ are daily changes in the ratio of traded put contracts and call contracts. In Table 2.4, we present the regression results of short-term options (Panel A, T=1 month) and long-term options (Panel B, T=6 months).

Table 2.4: Regression Results (Risk-Neutral Moments)

Panel A: 1 month				
	Volatility			
	Coefficient	T-stat		
Constant	-0.00	-0.17		
Dependent $_{-1}$	-0.06**	-2.03		
Returns $_{-1}$	-0.01	0.66		
Flows $_{-1}$	-0.07***	-4.34		
Put/Call ratio $_{-1}$	0.00	0.57		
R^2	1.9%			
	Skewness		Kurtosis	
	Coefficient	T-stat	Coefficient	Coefficient
Constant	0.00	0.04	0.00	0.05
Dependent $_{-1}$	-0.46***	-17.53	-0.41***	-15.07
Returns $_{-1}$	0.64	1.24	0.31	0.31
Flows $_{-1}$	0.03	0.07	1.67	1.56
Volatility $_{-1}$	1.74*	1.74	-5.28**	-2.55
Put/Call ratio $_{-1}$	0.00	0.89	-0.00	-0.4
R^2	22.5%		15.9%	

Panel B: 6 months				
	Volatility			
	Coefficient	T-stat		
Constant	-0.00	-0.33		
Dependent $_{-1}$	-0.03	-0.93		
Returns $_{-1}$	-0.05***	-2.79		
Flows $_{-1}$	-0.04**	-2.05		
Put/Call ratio $_{-1}$	0.00*	1.66		
R^2	1.2%			
	Skewness		Kurtosis	
	Coefficient	T-stat	Coefficient	T-stat
Constant	0.00	0.01	-0.00	-0.01
Dependent $_{-1}$	-0.51***	-20.30	-0.51***	-19.47
Returns $_{-1}$	1.40**	2.21	6.20**	2.42
Flows $_{-1}$	-1.48*	-1.89	-4.16	-1.34
Volatility $_{-1}$	4.86***	4.24	14.64***	3.21
Put/Call ratio $_{-1}$	-0.00	-0.59	-0.01	-0.71
R^2	25.2%		23.9%	

Note: This table shows the regression results of time series regressions according to equations (12-14), where risk-neutral moments of S&P500 index options are regressed on lagged explanatory variables. In panel A, we present the results for short-term options (T=1 month) and in panel B for long term options (T=6 months). The first row shows the dependent variable of each pair of regressions and the explanatory variables are listed in the first column. We report the respective estimate and t-value, where standard errors are adjusted for heteroscedasticity and serial correlation according to Newey and West (1987). *** indicates statistical significance at the 1 percent level, ** at the 5 percent level and * at the 10 percent level.

The dependent variables are the changes in risk-neutral volatility, skewness and kurtosis calculated model-free according to Equation (3). Independent variables are always the lagged dependent variable and other explanatory variables. The other explanatory variables enter the regression equation as lagged variables. Results suggest that for short-term options there is an insignificant relationship between previous day news from the stock market (proxied by the residuals from the return regressions) and changes in volatility. We see this as further evidence that all information in the stock market is immediately incorporated in option prices, resulting in a strong contemporaneous relationship between returns and implied volatilities that is often documented in previous research. On the other hand, the lagged relationship between noise (proxied by the residuals from the flow regressions) and changes in volatility is highly significant. In other words, if mutual fund / retail investors become bearish (bullish) on one day, risk-neutral volatilities of index options tend to increase (decrease) on the following day. The impact on implied volatilities is stronger for short-term options compared to long-term options.

Furthermore, in line with rational option pricing models and the results of Han (2008), an increase in the volatility on one day tends to reduce the negative skewness in the risk-neutral distribution, indicated by a positive coefficient. The results for kurtosis are mixed. While the effect of volatility is positive for long term options, for short term options an increase of volatility leads to a reduction of risk-neutral kurtosis on the following day. Interestingly, and in contrast to Han (2008), for short-term options, the lagged residuals of returns and flows, news and noise, do not affect the risk-neutral skewness and kurtosis.

In contrast, for longer-term options, the return innovations positively affect skewness and kurtosis. In other words, positive (negative) news from the stock market causes the risk-neutral skewness of long-term options to be less (more) negative on the following day, while it does not affect option prices of short-term options. The other control variable, the changes in put/call ratios, does not seem to explain the changes in risk-neutral skewness and kurtosis.

Therefore, results are partly in contrast to our findings using more ad-hoc measures of skewness and kurtosis. Based on those findings, we believe strongly that practical measures like risk-reversals and butterfly spreads (based on point estimates of the implied volatility surface) provide useful information about the shape of the implied volatility smile, but should not be considered to be proxies for skewness and kurtosis of the underlying risk-neutral distribution, which are based on all traded options over the whole delta range.

2.4 Conclusions

While there is substantial literature that looks at the impact of investor sentiment on stock prices, only a few studies investigate how option prices are affected by behavioral biases. Previous research suggests that investor sentiment is an important determinant of option prices and that the observed mispricing is sentiment-driven. The implied volatility surface is affected and, in particular, the index option smile is steeper (flatter) when the market is bearish (bullish). In this paper, using daily aggregated US equity fund flows as a proxy for noise trader sentiment, we investigate empirically if any noise trader activity affects the cross-section of S&P 500 index option prices. Other commonly used implicit or explicit measures of sentiment are typically only available at lower frequencies. Relying on those low frequency measures, previous findings suggest that market sentiment correlates contemporaneously with implied volatilities and the risk-neutral skewness, where the causality is unclear. Daily data allow us to investigate if the cross-section of index option prices is conditional on our beginning-of-period proxy of noise trading in the stock market and test for information as well as noise transmission effects.

Overall, our results suggest strongly that noise from the stock market is transmitted into the index option market. In particular, we find that the lagged relationship between noise (proxied by the residuals from the flow regressions) and changes in volatility is substantial and highly significant. We find that when noise traders are bearish (bullish) on a particular day, resulting in flows out of (into) US equity funds, volatilities of S&P500 index options tend to increase (decrease) significantly the following day. In line with other evidence, the effects are typically more pronounced for out-of-the-money options and short-term options. Furthermore, our findings also suggest that the shape of the option smirk is partly caused by noise traders active in the equity market. We find evidence that noise affects the shape of the smile, represented by ad-hoc measures that characterize the cross-section of index options like risk reversals or butterfly spreads, but not risk-neutral moments like skewness or kurtosis. We conclude that ad-hoc measures provide useful information about the shape of the implied volatility smile, but should not be considered to be proxies for skewness and kurtosis of the underlying risk-neutral distribution. Our results have implications for risk management and derivative pricing. More specifically, they provide guidance for options traders to favor the risk-neutral measures instead of the ad-hoc approach.

Chapter Three

Stein's Overreaction Puzzle: Anomaly or Perfectly Rational Behavior?[ii]

3.1 Introduction

This chapter deals with a phenomenon first highlighted by Stein (1989). In the options market, empirical studies have shown that longer-maturity options tend to overreact to changes in implied volatility of shorter-maturity options, while finance theory predicts a smoothing effect on a time series of implied volatilities with the same underlying asset but different maturities. This empirical contradiction of theoretical expectations creates a puzzle.

Stein (1989) derives and tests empirically a model that describes the relationship between implied volatilities on options of different maturities. Assuming that volatility evolves according to a given continuous-time mean-reverting AR1 process (with a constant long-run mean and a constant coefficient of mean-reversion) he finds that, theoretically, the implied volatility of longer maturity (two-month) options should move in a responsive, but smoothing manner to changes in implied volatility of shorter maturity (one-month) options. However, the empirical values of this elasticity exceeded the theoretical upper bound of normal-reaction. He interprets his findings as overreaction, (caused by market inefficiencies) claiming that this contradicts the rational expectations hypothesis for the term structure of implied volatilities.

Poteshman (2001) extends the work of Stein (1989), and constructs two variables from index options, namely two risk-neutral instantaneous variances from nearby ATM call and ATM put options and from distant ATM call and ATM put options, respectively. By detecting that the difference between the former variance implied from long-maturity options and the latter variance implied from short-maturity options is increasing in the level of an instantaneous variance which minimizes the pricing errors under Heston (1993) model, the author states that long-horizon overreaction is present.

Those 'term structure' tests suggest that implied volatilities of long-term options react quite strongly to changes in implied volatilities of short-term options. Phrasing it differently, long-term options do not seem to display the rationally expected smoothing behavior. Given the observed level of mean-reversion in volatility, Stein (1989) and Poteshman (2001) interpreted those findings as evidence for overreaction in the options market. In this study, we challeng this view.

In line with those 'term structure' tests, Christoffersen et al. (2013), replicate Stein's analysis with more recent data (1996 - 2009). Considering the same maturity time frame[6], they demonstrate the robustness of Stein's results but differ in interpreting them. While Stein considers the overreaction observed in his sample as an anomaly vis-à-vis rational expectations, Christoffersen et al. (2013) explain it by a variance depended pricing kernel[7]. Our study builds on their theoretical results.

Assuming a stochastic variance process in a rational expectation framework, we show theoretically that under normal market conditions the risk-neutral volatility process is substantially more persistent than the physical one. Investors' risk aversion appears to be the main factor driving this persistence. Theoretically, long-term volatility should react more strongly to changes in short-term volatility in periods when investors are highly risk averse, and risk-neutral volatility is highly persistent. In contrast, in periods of low risk aversion, long-term volatility should react less strongly to changes in short-term volatility,

[6] One-month maturity for short-term options and two-month maturity for long-term option
[7] In their model, the pricing kernel specification is monotonic in returns and also monotonic in volatility.

because risk-neutral volatility is less persistent. Using daily data on S&P 500 index options for the 2000-2010 period, we verify empirically these theoretical predictions. In periods of high risk aversion, long-term volatility reacts strongly to changes in short-term volatility, which can be explained by highly persistent risk-neutral volatility dynamics in that period. The effect cannot be observed in periods of low risk aversion, because of less persistent volatility dynamics. Overall, we provide strong evidence that the empirical observation that Stein (1989) discovered is not overreaction, but in line with perfectly rational behavior.

The remainder of this chapter is organized as follows. In the next section, we present the theoretical framework, in Section 3.3 and 3.4, we discuss the data and present the empirical analysis and results and Section 3.5 concludes.

3.2 Theoretical Motivation

We illustrate the relationship between risk aversion, risk-neutral mean-reversion in volatility and Stein's overreaction hypothesis using a stochastic variance process in a rational expectation framework. In theory, we show that in a market characterized by highly risk-averse investors, we expect to observe a highly persistent risk-neutral variance process. The stylized fact that Stein (1989) discovered and interpreted as "overreaction" is not an anomaly of the option market, but can be shown to be perfectly in line with a rational expectations framework.

A typical stochastic volatility framework for the dynamics of the spot price is Heston (1993) model

$$dS(t) = (r + \mu v(t))S(t)dt + \sqrt{v(t)}S(t)dZ_1(t) \tag{3.1}$$

$$dv(t) = \kappa[\theta - v(t)]dt + \sigma\sqrt{v(t)}dZ_2(\mathrm{t}) \tag{3.2}$$

where $S(t)$ is the spot price; $v(t)$ is instantaneous variance; r is risk-free rate; μ relates to the equity risk premium; Wiener process $Z_2(t)$ has a correlation ρ with $Z_1(t)$. In a risk-neutral world, the variance follows the same mean reverting process but under risk-neutral measures

$$dS(t) = rS(t)dt + \sqrt{v(t)}S(t)dZ_1^*(t) \tag{3.3}$$

$$dv(t) = \kappa^*[\theta^* - v(t)]dt + \sigma\sqrt{v(t)}dZ_2^*(t) \tag{3.4}$$

$$\text{where } \kappa^* = \kappa + \lambda \text{ and } \theta^* = \kappa\theta/(\kappa + \lambda) \tag{3.5}$$

Therefore, under risk-neutral pricing probabilities, the variance drifts towards a long run mean θ^*, with a mean reverting speed κ^*. The unspecified term λ is the variance risk premium and is mathematically a compensation transferred to the drift term from the change of probability measures.

Furthermore, we assume that the representative investor adopts a CRRA utility function as

$$U(c) = \begin{cases} \dfrac{c^{1-\gamma}}{1-\gamma}, & \gamma > 0, \gamma \neq 1 \\ \ln c, & \gamma = 1 \end{cases}$$

Without loss of generality, we assume a two-period model. By maximizing utility over the two periods when investors choose to either allocate their wealth on stocks or to consume, in equilibrium, we get a stock price process satisfying a martingale pricing condition as $\pi_t S_t = E_t(\pi_T S_T)$. The pricing kernel, based on consumption return state (c_T/c_t) is thus, as follows:

$$\frac{\pi_T}{\pi_t} = \beta\left(\frac{c_T}{c_t}\right)^{-\gamma}$$

Where the coefficient β is time preference function; $\gamma = \frac{-cU``(c)}{U`(c)}$ is a measure of the relative risk aversion.

Out of this general approach, instead of using consumption data, we follow empirical pricing kernel research and replace the consumption return, (c_T/c_t), with a proxy for the stock return (S_T/S_t).[8]

Depending on where Taylor expansion series is truncated of the stock price process, pricing kernel take forms of different degree of complications. We follow Christoffersen et al.'s (2013) specification of the pricing kernel for the Heston stock price process as follows:

$$\frac{\pi_T}{\pi_t} = \left(\frac{S_T}{S_t}\right)^{-\gamma} \exp\left(\delta(T-t)T + \eta \int_t^T v(s)ds + \xi(v(T) - v(t))\right) \tag{3.6}$$

where δ and η determine investor's return preference; ξ determines the variance preference; and γ is constant relative risk aversion coefficient. Compared with the basic pricing kernel, $\frac{\pi_T}{\pi_t} = \left(\frac{S_T}{S_t}\right)^{-\gamma}$, the more general (3.6) takes accounts of not only risk aversion, but also investor's variance preference. Out of arbitrage, the variance preference ξ is positive as hedging needs increase when we expect the pricing kernel to be increasing in the variance.

By applying the martingale process with the above pricing kernel, the variance risk premium is derived as in Christoffersen et al (2013)

$$\lambda = \rho\sigma\gamma - \sigma^2\xi \tag{3.7}$$

where the correlation between a stock market return and variance is empirically found to be negative, $\rho < 0$.

[8] The disadvantage of using consumption data is that there are measurement problems and are not suited to identify time variation in pricing kernel parameters. See for example Jackwerth (2000), Ait-Sahalia and Lo (2000).

Our focus is on the components and empirical implications of the variance risk premium. Among papers that study anomalies in options market, [see Ait-Sahalia and Lo (2000); Shive and Shumway (2009); Bakshi, Madan and Panayotov (2010); Christoffersen, Heston and Jacobs (2013); Lehnert, Lin and Wolff (2013); etc.], it is unanimously found that there is on average a negative variance risk premium, $\lambda < 0$, be it from a pricing kernel based on kinds of varieties of stochastic models or from a 'by-product' of general equilibrium models.

From (3.7), we note that the variance risk premium λ is closely related to investor's risk aversion γ and variance preference ξ. When investors are risk averse, $\gamma > 0$, and, therefore ξ being positive[9], investors require more compensation for risk, both terms in (3.7) will be negative, which creates a negative variance risk premium. Furthermore, the higher the risk aversion, the more negative the variance risk premium is, the larger the difference between the mean-reversion of the physical volatility process and the mean-reversion of the risk-neutral process. Once there is a reasonably high risk aversion, it would lead to a lower rate of mean reversion, or a more persistent risk-neutral volatility process, as in $\kappa^* = \kappa + \lambda = k + (\rho\sigma\gamma - \sigma^2\xi) < \kappa$.

In a related study, Lehnert et al (2013) solve for the equity risk premium in a general equilibrium framework with a CRRA representative investor. They find that the equilibrium risk premium is a function greatly determined by a representative investor's risk aversion, which is found to be time-varying. In their empirical analysis, they show that the time-variation in investor sentiment can be associated with time-varying risk aversion. During times of low investor sentiment, risk aversion is high; when noise traders demand for equity increases and sentiment is high, risk aversion decreases significantly. Therefore, in periods of low sentiment, investors are more risk averse, which, according to the above theoretical finding, leads to a more persistent risk-neutral volatility process. In those periods,

[9] In Christoffersen et al. (2013), it is argued that out of arbitrage, the variance preference is positive as hedging needs increase when the pricing kernel is expected to increase in the variance.

we would expect to find the stylized fact that Stein (1989) interpreted as "overreaction" to be significant. However, it would not be an option anomaly, but fully consistent with a rational expectation framework. On the other hand, according to the empirical findings of Lehnert et al. (2013), high sentiment periods correspond to periods of low investor risk aversion and, therefore, a lower variance preference. In those periods, according to the relationship, $\kappa^* = \kappa + \lambda = k + (\rho\sigma\gamma - \sigma^2\xi)$, the risk-neutral mean-reversion is theoretically stronger. As a result, during high sentiment periods, we expect the stylized fact that Stein (1989) discovered, not to be a feature of the data. In the next section, we proceed by testing our hypothesis using S&P500 index options.

3.3 Data

We use daily European option data on the S&P 500 index (symbol: SPX) from OptionMetrics over the period January 2000 to April 2010. We follow Barone-Adesi et al. (2008) in filtering the original options data. We focus on all traded options with more than a week, but less than one year maturity. For liquidity reasons, we only consider closing prices of out-of-the money put and call SPX options for each trading day. The bid-ask midpoint prices are taken. In line with Bollen and Whaley (2004), we exclude options with absolute call deltas below 0.02 or above 0.98 because of distortions caused by price discreteness. The underlying S&P 500 index level, dividend yields and the term structure of zero-coupon default-free interest rates are also provided by OptionMetrics. On each day we fit a functional form with curvature to the term structure in order to obtain the interest rate that matches the maturity of the option. We price the options by using the dividend-adjusted underlying S&P level.

For our empirical analysis, we use a modification of the ad-hoc Black-Scholes model of Dumas, Fleming and Whaley (1998) to estimate the implied volatility surface of index options. The aim is to use all available information content in index option prices and to investigate the time series of standardized implied volatilities for fixed moneyness options with fixed time to maturities. Rather than averaging the two contracts that are closest to at-the-money or closest to one-month maturity, we fit a modified ad-hoc Black-Scholes model to all option contracts on a given day. Subsequently, we obtain the desired implied volatility and option prices that correspond to a particular moneyness and maturity. This strategy eliminates successfully some of the noise from the data (see Christoffersen et al. (2013)).

As indicated in Bollen and Whaley (2004), it is industry practice to quote Black-Scholes volatilities by option delta. Therefore, we allow each option to have its own Black-Scholes implied volatility depending on the options delta and time to maturity T. We use the following functional form for the options implied volatility[10]:

$$IV_{i,j} = \alpha_0 + \alpha_1 \text{delta}_{C\,i,j} + \alpha_2 \text{delta}_{C\,i,j}^{2} + \alpha_3 T_j + \alpha_4 T_j^2 + \alpha_5 \text{delta}_{C\,i,j} T_j \,,$$

where IV_{ij} denotes the Black-Scholes implied volatility and $\text{delta}_{C,i,j}$, the call delta of an option for the i-th exercise price and j-th maturity[11]. T_j denotes the time to maturity of an option for the j-th maturity.

In order to test for "overreaction" in our sample, we follow precisely the methodology of Stein (1989) and Christoffersen et al. (2013). But while they use implied at-the-money volatilities, we use a model-free method applied to option prices to obtain the variance of the risk-neutral distribution (Bakshi et al. (2003))[12]. In recent years, the approach became very popular in the empirical literature studying option markets (see e.g. Han (2008) for index options and Bekkour et al. (2013) for exchange rate options). We derive model-free measures of risk-neutral variance (*$Var_t(T)$)* based on put and call option prices obtained over the complete moneyness range for a particular time to maturity T,

$$\text{Var}_t(T) = e^{rT} V_t(T) - \mu_t^2(T)$$

$$\mu_t(T) = e^{rT} - 1 - \frac{e^{rT}}{2} V_t(T) - \frac{e^{rT}}{6} W_t(T) - \frac{e^{rT}}{24} X_t(T)$$

$$\text{V}_t(T) = \int_{S_t}^{\infty} \frac{2(1 - \ln(\frac{K}{S_t}))}{K^2} C_t(T,K)dK + \int_0^{S_t} \frac{2(1 + \ln(\frac{S_t}{K}))}{K^2} P_t(T,K)dK$$

[10] We have also tried other specifications for the functional form that are frequently used in the literature (replacing delta by exercise price K or moneyness K/F, where F is the forward rate, but the results are robust to these changes).
[11] For put options, we use the corresponding call delta in the implied volatility regression.
[12] In the empirical part of the paper, we replicate the analysis with implied volatilities obtained by interpolating near-the-money short-term options or obtained using a functional form for the implied volatility surface, and the results are consistent.

$$W_t(T) = \int_{S_t}^{\infty} \frac{6\ln(\frac{K}{S_t})) - 3\left[\ln(\frac{K}{S_t})\right]^2}{K^2} C_t(T,K)dK - \int_0^{S_t} \frac{6\ln(\frac{S_t}{K})) - 3\left[\ln(\frac{S_t}{K})\right]^2}{K^2} P_t(T,K)dK$$

$$X_t(T) = \int_{S_t}^{\infty} \frac{12\ln(\frac{K}{S_t})) - 4\left[\ln(\frac{K}{S_t})\right]^3}{K^2} C_t(T,K)dK + \int_0^{S_t} \frac{12\ln(\frac{S_t}{K})) - 4\left[\ln(\frac{S_t}{K})\right]^3}{K^2} P_t(T,K)dK$$

With C being the price of a call option and P, the price of a put option. S is the dividend adjusted index level, K is the option strike price, T is the time to maturity and r is the risk-free rate. All risk-neutral variances ($Var_t(T)$) corresponding to a time to maturity T are transformed into the annualized risk-neutral volatilities ($Vol_t(T)$) which are used in the subsequent analysis.

In line with Stein (1989) and Christoffersen et al. (2013), we consider one month to be short-term and two months to be long-term. Additionally, as a robustness check, we replicate the analysis with three, six and nine months, that we consider to be the longer term. Table I presents an overview of risk-neutral annualized volatilities for short as well as long-term options.

Table 3.1
Options Volatility - Summary Statistics
Summary statistics of options' risk-neutral volatilities.

Maturity (months)	Mean	Standard Deviation	Min	Max	AR 1	N
Panel A: Full Sample						
1	0.206	0.097	0.088	0.817	0.981	2581
2	0.213	0.094	0.095	0.807	0.986	2581
3	0.219	0.091	0.101	0.777	0.990	2581
6	0.230	0.086	0.114	0.662	0.995	2581
9	0.233	0.082	0.120	0.603	0.993	2581

Panel B: Low Sentiment Period						
1	0.206	0.102	0.088	0.672	0.981	1120
2	0.216	0.103	0.095	0.667	0.984	1120
3	0.225	0.103	0.101	0.657	0.986	1120
6	0.242	0.104	0.114	0.614	0.989	1120
9	0.245	0.100	0.120	0.570	0.988	1120
Panel C: High Sentiment Period						
1	0.207	0.093	0.088	0.817	0.971	1440
2	0.211	0.087	0.096	0.807	0.977	1440
3	0.215	0.081	0.104	0.777	0.982	1440
6	0.221	0.069	0.119	0.662	0.989	1440
9	0.223	0.065	0.124	0.603	0.987	1440

The average volatility (annualized risk-neutral variance) over the entire sample period and for all maturities is 0.22, which does not come as a surprise as our sample encompasses periods of both low volatility (post-dotcom bubble years) and greater volatility (2007-2009 crisis). On average the term structure of risk-neutral volatility is upward sloping, where short-term volatilities tend to fluctuate more than long-term volatilities. As can be seen from the maximum figures, in particular during periods of market stress, short-term volatilities can increase by substantially more than long-term volatilities. Once we subdivide the sample period into periods of high and low sentiment[13], we find that average volatilities are quite similar, but in high sentiment periods, volatilities can increase substantially, as can be seen from the maximum values. Additionally, estimates of first-order autocorrelation of the volatility time series suggest that in low sentiment periods the risk-neutral volatility process is more persistent compared to high sentiment periods: a preliminary finding that supports our main hypothesis. In the next section, we proceed with the term-structure tests.

[13] As in Lehnert et al. (2013), the average of the past six months' Baker and Wurgler (2006; 2007) end-of-month sentiment index is considered to be the current-month index. We thank Malcolm Baker and Jeffrey Wurgler for making the data available.

3.4 Empirical Testing

In order to test if investors are able to incorporate new information correctly into option prices, Stein (1989) derives an elasticity relationship using two options on the same underlying asset but with different time-to-maturity: a short-term maturity option, with time-to-maturity of e.g. one month and annualized volatility Vol^{st}_{t}, and a long-term maturity option, with time-to-maturity of e.g. two months and annualized volatility Vol^{lt}_{t}.

The elasticity relationship may be expressed as[14]:

$$\left(Vol_t^{lt} - iV\right) = ½\left(Vol_t^{st} - iV\right)E_t\left(Vol_{t+(lt-st)}^{st} - iV\right)$$

Where iV is the instantaneous volatility. It can be rearranged into:

$$E_t\left[\left(Vol_{t+(lt-st)}^{st} - Vol_t^{st}\right) - 2\left(Vol_t^{lt} - Vol_t^{st}\right)\right] = 0$$

Motivated by empirical evidence of mean-reversion in volatility, Stein (1989) hypothesized that under 'normal' reaction, the prediction error $E_t[.]$ should remain white noise. In the case of what he considered to be "overreaction", the prediction error will be negatively correlated with Vol^{st}_{t}. For the same reason, a positive correlation of the prediction error with Vol^{st}_{t} could suggest an "underreaction" phenomenon[15]. We follow Stein (1989) and Christoffersen et al. (2013) and regress the prediction error defined earlier on the short-term volatility. We implement the regression approach using the daily time series of one month and two month risk-neutral volatilities[16].

$$(Vol_{t+21}^{1m} - Vol_t^{1m}) - 2(Vol_t^{2m} - Vol_t^{1m}) = \alpha + \beta Vol_t^{1m} + \varepsilon_{t+21}$$

---

[14] See Stein (1989) for details and full derivation.

[15] According to his tests, 'normal reaction' is supposed to yield insignificant results when the prediction error is regressed on Vol^{st}.

[16] We have also conducted the analysis with weekly data as in Stein (1989) and Christoffersen et al. (2013), results are robust to this change in frequency.

Here, short-term options are assumed to be one-month options and long-term options are assumed to be two-months options. The difference between the two terms is assumed to be one month or 21 trading days. All regressions are standard OLS and the results are displayed in Table 3.2.

Table 3.2
Prediction error against short-term risk-neutral volatility

$$(Vol_{t+21}^{1m} - Vol_t^{1m}) - 2(Vol_t^{2m} - Vol_t^{1m}) = \alpha + \beta Vol_t^{1m} + \varepsilon_{t+21}$$

Vol_t^{1m} is the risk-neutral volatility of options with short-term maturity (1 month). Vol_t^{2m} is the risk-neutral volatility of options with long-term maturity (2 months).

Sample Period	Regression Coefficient	Standard Error	t-Statistic	N
2000	-0.295	0.061	-4.82	252
2001	-0.513	0.065	-7.89	248
2002	-0.238	0.047	-5.01	252
2003	-0.103	0.032	-3.25	252
2004	-0.617	0.057	-10.87	251
2005	-0.664	0.063	-10.48	252
2006	-0.264	0.064	-4.12	251
2007	-0.354	0.048	-7.37	251
2008	-0.180	0.051	-3.56	250
2009	-0.089	0.024	-3.70	252
Full Sample	-0.099	0.012	-8.68	2560

Our results are in line with common findings in the literature. The estimates displayed in Table 3.2 are indeed consistent in magnitude with the previously cited works. Ranging from -0.09 to -0.66, all the regression coefficients of the year-by-year analysis in our sample are significantly negative. As in the existing literature, the regression coefficient of the full sample falls into the lower bound of the range cited above (-0.1) but still remains significantly different to what Stein (1989) would consider to be "normal" reaction. However, interestingly, the coefficients presented in Table 3.2, as well as the findings reported in Stein (1989) and Christoffersen et al. (2013), suggest that there is quite some variation over time. In order to further investigate this observation, we perform rolling-window semester regressions. Figure I shows the estimated regression coefficients together with the lower and upper bounds of the 99% confidence interval. It is apparent from the graph that most of the time and in line with previous evidence, the coefficient is significantly negative. However, for certain periods, the analysis leads to insignificant results, which suggests that the effect depends on other factors. Therefore, the outcome of the term-structure tests that we perform might depend on particular market conditions; a claim that we motivated theoretically in section 3.2, which we would like to test empirically in the remaining part of this chapter.

Figure 3.1
Rolling-Window Semester Regressions

$$(Vol_{t+21}^{1m} - Vol_t^{1m}) - 2(Vol_t^{2m} - Vol_t^{1m}) = \alpha + \beta Vol_t^{1m} + \varepsilon_{t+21}$$

Vol_t^{1m} is the risk-neutral volatility of options with short-term maturity (one month). Vol_t^{2m} is the risk-neutral volatility of options with long-term maturity (two months). The figure shows the estimated regression coefficients together with the lower and upper bounds of the 99% confidence interval. The rolling-window regressions are performed on a daily data basis over the previous six months using 126 daily observations. For example, the first data point refers to values of the estimated coefficient, where the regression is performed over the first half of the year 2000.

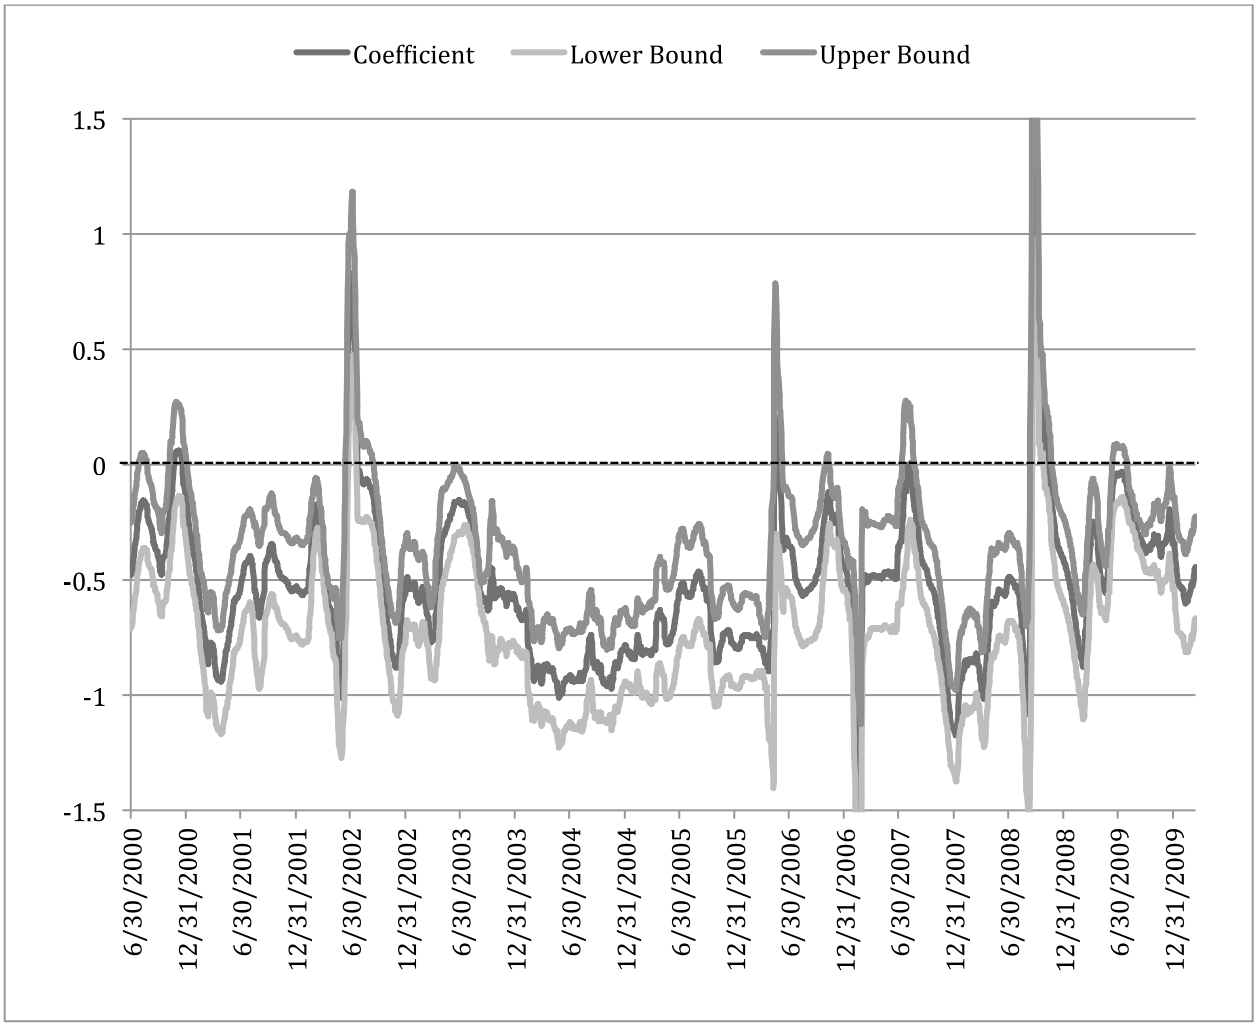

As motivated in the theoretical part of the paper, we hypothesize that the previous findings can be explained by a more persistent risk-neutral volatility process and, therefore, should not be interpreted as an option anomaly. Additionally, given that the degree of persistency depends on investors' risk aversion, it would be interesting to investigate the robustness of the argument. Hence, following our theoretical reasoning and the empirical results of Lehnert et al. (2013), we test the relationship between the prediction error and short-term volatility under different market conditions. As in Lehnert et al. (2013), the average of the past six months' Baker and Wurgler (2006; 2007) end-of-previous-month sentiment index is considered to be the current-month index. This method allows us to smooth out some noise in the data[17]. An observation is regarded as in a low sentiment regime (high risk aversion) if our sentiment index is below zero and as in a high sentiment regime (low risk aversion) if it is above zero. Subsequently, we regress the prediction error on the short-term volatility, but control for the impact of time variation in investors risk aversion, proxied by sentiment[18]. In particular, we run the following regression

$$(Vol_{t+21}^{1m} - Vol_t^{1m}) - 2(Vol_t^{2m} - Vol_t^{1m}) = \alpha + \alpha_H D_H + \beta Vol_t^{1m} + \beta_H D_H Vol_t^{1m} + \varepsilon_{t+21}$$

where the main variables are the same as previously, but D_H is a dummy that is one during high sentiment (low risk aversion) periods and zero otherwise. Results for the whole period are presented in Table III.

[17] We have also constructed other sentiment indicators using the original Baker and Wurgler (2006; 2007) data, but there is not qualitative change in the results.
[18] See Lehnert et al. (2013) for details.

Table 3.3
Prediction error against short-term risk-neutral volatility with sentiment

$$(Vol_{t+21}^{1m} - Vol_t^{1m}) - 2(Vol_t^{2m} - Vol_t^{1m}) = \alpha + \alpha_H D_H + \beta Vol_t^{1m} + \beta_H D_H Vol_t^{1m} + \varepsilon_{t+21}$$

Vol_t^{1m} is the risk-neutral volatility of options with short-term maturity (one month). Vol_t^{2m} is the risk-neutral volatility of options with long-term maturity (two months). D_H is a dummy variable that is equal to 1 during high sentiment periods and 0 otherwise.

	Coefficient	Standard Error	t-Statistic	N
α	0.014	0.004	3.78	2560
α_H	-0.016	0.005	-3.13	2560
β	-0.205	0.016	-12.98	2560
β_H	0.203	0.022	9.24	2560

Overall, the empirical results support the theoretical prediction that the relationship between the prediction error and short-term volatility varies with investors' risk aversion. In periods of high risk aversion (proxied by sentiment being low) the relationship is highly significant (β=-0.205 with a t-statistic of -12.98) and in periods of low risk aversion (proxied by sentiment being high) the relationship is weakened dramatically (β_H=0.203 with a t-statistic of 9.24). As a result, in the high sentiment period, the relationship is essentially flat ($\beta + \beta_H$=-0.002). In addition, the two-regime regression accommodates the data much better than the one-regime equation, with R^2 increasing from less than 3% to more than 11%.

Overall, in line with our theoretical motivation, risk aversion and, therefore, the degree of persistence of the risk-neutral volatility process explains the strong reaction of long-term volatilities to changes in short-term volatility, the empirical observation that Stein (1989) discovered. It should not be interpreted as "overreaction", because it is in line with perfectly rational behavior. The absence of the relationship in periods that can be characterized by low risk aversion, strengthens substantially the argument in favor of a rational explanation.

As a robustness check, we ran the same regressions over the complete term structure of risk-neutral volatility. Long-term options are considered to be three, six, and nine months, while the short-term is considered to be one month. This examination also represents a robustness check of the Christoffersen et al. (2013) and Stein (1989) analysis. The results are presented in Tables 3.4 and 3.5.

Table 3.4
Prediction error against short-term risk-neutral volatility

$$(Vol_{t+X}^{1m} - Vol_t^{1m}) - 2(Vol_t^{lm} - Vol_t^{1m}) = \alpha + \beta Vol_t^{1m} + \varepsilon_{t+X}$$

Vol_t^{1m} is the risk-neutral volatility of options with short-term maturity (fixed at 1 month). Vol_t^{Km} is the risk-neutral volatility of options with long-term maturity (l=3, 6 and 9 months). X corresponds to the time difference (in trading days) between long-term options and short-term options.

Panel A : 3 months				
Sample Period	Sample Period	Sample Period	Sample Period	Sample Period
2000	2000	2000	2000	2000
2001	2001	2001	2001	2001
2002	2002	2002	2002	2002
2003	2003	2003	2003	2003
2004	2004	2004	2004	2004
2005	2005	2005	2005	2005
2006	2006	2006	2006	2006
2007	2007	2007	2007	2007
2008	2008	2008	2008	2008
2009	2009	2009	2009	2009
Full Sample	Full Sample	Full Sample	Full Sample	Full Sample
Panel A : 6 months				
Sample Period	Sample Period	Sample Period	Sample Period	Sample Period
2000	2000	2000	2000	2000
2001	2001	2001	2001	2001
2002	2002	2002	2002	2002
2003	2003	2003	2003	2003
2004	2004	2004	2004	2004
2005	2005	2005	2005	2005
2006	2006	2006	2006	2006
2007	2007	2007	2007	2007
2008	2008	2008	2008	2008
2009	2009	2009	2009	2009
Full Sample	Full Sample	Full Sample	Full Sample	Full Sample
Panel A : 9 months				
Sample Period	Sample Period	Sample Period	Sample Period	Sample Period
2000	2000	2000	2000	2000
2001	2001	2001	2001	2001
2002	2002	2002	2002	2002
2003	2003	2003	2003	2003
2004	2004	2004	2004	2004
2005	2005	2005	2005	2005
2006	2006	2006	2006	2006
2007	2007	2007	2007	2007
2008	2008	2008	2008	2008
2009	2009	2009	2009	2009
Full Sample	Full Sample	Full Sample	Full Sample	Full Sample

Table 3.5
Prediction error against short-term risk-neutral volatility with sentiment

$$(Vol_{t+X}^{1m} - Vol_t^{1m}) - 2(Vol_t^{lm} - Vol_t^{1m}) = \alpha + \alpha_H D_H + \beta Vol_t^{1m} + \beta_H D_H Vol_t^{1m} + \varepsilon_{t+X}$$

Vol_t^{1m} is the risk-neutral volatility of options with short-term maturity (fixed at 1 month). Vol_t^{Km} is the risk-neutral volatility of options with long-term maturity (l=3, 6 and 9 months). X corresponds to the time difference (in trading days) between long-term options and short-term options. D_H is a dummy variable that is equal to 1 during high sentiment periods and 0 otherwise.

	Coefficient	Standard Error	t-Statistic	N
Panel A : 3 months				
α	0.012	0.005	2.56	2539
α_H	-0.008	0.007	-1.23	2539
β	-0.318	0.021	-15.32	2539
β_H	0.283	0.029	9.79	2539
Panel B : 6 months				
	Coefficient	Standard Error	t-Statistic	N
α	0.003	0.006	0.51	2476
α_H	-0.017	0.008	-2.08	2476
β	-0.534	0.025	-20.97	2476
β_H	0.594	0.035	16.83	2476
Panel C : 9 months				
	Coefficient	Standard Error	t-Statistic	N
α	0.018	0.005	3.81	2413
α_H	-0.007	0.006	-1.12	2413
β	-0.477	0.020	-23.81	2413
β_H	0.520	0.028	18.77	2413

Results suggest that not only do all our previous findings hold, but they become even stronger once we consider different parts of the term structure. All coefficients are significantly negative, except for the year 2000 for the six and nine month options. Once we control for the impact of time variation in investors' risk aversion (Table V), regression results show the same pattern that we observed in the previous analysis, which were remarkably stable over the complete term structure.

3.5 Conclusion

The findings of Stein (1989) suggest that implied volatilities of long-term options react “too strongly” to changes in implied volatilities of short-term options and do not display the rationally expected smoothing behavior. Given the observed strong mean-reversion in volatility, Stein (1989) interpreted the results as evidence for overreaction in the options market, but Christoffersen et al. (2013) and our findings challenge this view. Building on a stochastic variance process in a rational expectation framework, we show theoretically that under normal market conditions the risk-neutral volatility process is substantially more persistent than the physical one. Investors’ risk aversion appears to be the main factor driving this persistence. Theoretically, long-term volatility should react more strongly to changes in short-term volatility in periods when investors are highly risk averse, and risk-neutral volatility is highly persistent. In contrast, in periods of low risk aversion, long-term volatility should react less strongly to changes in short-term volatility, because risk-neutral volatility is less persistent. Using daily data on S&P 500 index options for the 2000-2010 period, we verify empirically these theoretical predictions. In periods of high risk aversion, long-term volatility react strongly to changes in short-term volatility, which can be explained by the high persistence of risk-neutral volatility in that period. The effect cannot be observed in periods of low risk aversion, because of a less persistent volatility process. Overall, we provide strong evidence that the empirical observation that Stein (1989) discovered is not overreaction, but in line with perfectly rational behavior.

For practitioners, our results strongly suggest the use of a risk-dependent approach when it comes to pricing equity options with different maturities.

Chapter Four

Is There a Bubble in the Art Market?[iii]

4.1 Introduction

Is there a bubble in the art market? This question was raised on November 15, 2013, by CNN journalist Ben Rooney, after a record-setting week for the art market that broke the highest price ever paid for an auctioned painting (Francis Bacon's 1969 triptych portrait of Lucian Freud fetched $142.4 million on November 12, 2013 at Christie's New York). The same night, the highest price was paid for the work of a living artist ($58.4 million for Jeff Koon's 1994 Balloon Dog – Orange, also auctioned at Christie's). These contributed to the auction house recording its most valuable art sale ever, with all works being sold for a total of more than $691 million. Earlier that year, hedge fund manager Steven Cohen had purchased Picasso's Le Rêve (1932) for about $155 million on a private deal. This price was only topped by one of the five versions of Paul Cezanne's The Card Players (1892-1893) bought privately by the State of Qatar in May 2011 for a reported price of more than $250 million. New York Times journalist Conrad de Aenlle recently argued that "[...] wealthy investors have bid up prices of works by a handful of Contemporary artists so high that it's turning the heads of collectors of more modest means". He continues by speculating that "[...] such a concentration of interest and money in one segment of a market leads to instability down the road [...]", and "[...] authorities on art investment worry that the budding romance could end quickly and badly".

But such a large rise in the prices for works of art does not necessarily means that there is a speculative bubble in the art market. Bubbles are generally defined as high-volume trading of a given category of assets at prices that are far above their intrinsic value (King et al. (1993)).

Given this definition, one can isolate the following constraints to the detection of bubbles: (i) high volumes must be traded, and (ii) prices must be above their fundamental values. When it comes to the art market, those constraints raise two questions:

- First, does the recent dramatic increase in art prices concern the entire art market or is it a "top one percent" phenomenon that benefits only to a few "brand-name" artists (high volumes constraint)?
- Second and most puzzling, how can we determine the intrinsic value of an artwork? If we cannot determine the intrinsic value it is impossible to detect a speculative bubble.

Stein (1977) notes that artwork and more particularly paintings are peculiar financial assets, being at the same time consumable goods and financial assets. This market began in the 16th century when collectors started to acquire works of art for aesthetic reasons, social status and also as an investment (Getty Museum (2013)). However, as Campbell (2008) notes, art markets still remain somewhat opaque and rather illiquid. Theyare further characterized by significant transaction costs and high barriers to entry (Kräussl and Wiehenkamp (2012)). Moreover, the value of artwork depends on a wide range of variables which are sometimes difficult to define, such as individual taste, fashion effects (Chanel 1995) or even the location of the sale.

In spite of all those specificities, economists like William Baumol had started to consider art as a financial asset as early as the 1960s and the interest for art as a financial asset has increased along with the growth of the art market itself. Beyond scholars and economists, the art market has also caught the interest of financial institutions (e.g. UBS, Credit Suisse, Deutsche Bank, BNP Paribas, Société Générale), which provide art advisory services to their wealthiest clients (Kräussl and Wiehenkamp, 2012).

As of 2012, the auction market for art represented a total of $12.3 billion worldwide (Artprice (2013)). In a broader perspective, McAndrew (2008) estimates the global art market (auctions and private deals) to be over $3 trillion with $50 billion annual turnover. Artprice (2013) reports that in 2013, 80% of the transactions on the global auction market involved works of art priced at or under $5,000. Only 1% of all auctioned paintings over the period 1970 to 2013 fetched prices higher than $1 million. Therefore, the high-end, record-level pieces discussed in the media only represent a tiny part of the global market. Taking this into consideration, global art markets numbers and the first question raised, we ran tests on different art market segments, namely Impressionist and Modern, Post-war and Contemporary, American, and Latin American. This enabled us to capture a broad spectrum of the art market.

Answering the second question is not that straightforward but it is key to understanding this topic. Even though Frey and Pommerehne (1989), Gérard-Varet (1995) and Chanel et al. (1996) list a few variables which are supposed to explain the formation of artwork prices (e.g. production cost, size and type of work, buyer income, measure of aesthetic quality, measures of artists' attributes, etc.), trying to determine a work of art's fundamental value is almost an impossible feat. Under rational expectations (i.e.the fundamental value of an asset being equal to its discounted expected stream of cash flows (present value theory)), it is relatively easy to obtain the expected cash flow earned by equities (dividend) or real estate (rent). The ownership of an artwork, on the other hand, provides no claim for monetary return other than some kind of convenience yield, which is also described as a "dividend of enjoyment" by Campbell (2008) and as "aesthetic pleasure" by Gérard-Varet (1995). However, this is so closely dependent on the motivations and characteristics of each individual owner that it makes it impossible to quantify it clearly. To overcome this 'fundamental value' issue, we use a new and direct method of bubble detection developed by Phillips et al. (2011). This approach is based on a right-tailed ADF (augmented Dickey-Fuller) test, which can detect directly any explosive behavior in a time series.

Our empirical findings suggest that there is strong evidence of a speculative bubble - in the mania phase of its formation - both in the Post-war and Contemporary and in the American art market segments. However, we do not find such a phenomenon in the other markets of our study, which makes the Contemporary and the American segments the most likely to develop bubble-type behaviors. The well-documented bubble of 1990 that features in our results serves as a solid robustness check for the methodology used.

The remaining part of this chapter is structured as follows: Section 4.2 motivates the theory behind bubble detection and reviews the different testing methods, Section 4.3 presents the data and the corresponding art market indices, Section 4.4 discusses our empirical findings and Section 4.5 concludes.

4.2 Framework

The desire to detect the formation of asset price bubbles has captured the attention of scholars and researchers over the last four decades. From the seminal work of Kindleberger (1978) to the recent modeling approach of Phillips et al. (2011), numerous statistical tests have been developed to identify the presence of bubbles in a time series. The basis for all approaches to this question is the definition of a bubble as: a sharp rise in a given asset price above a level sustainable by some fundamental values, followed by a sudden collapse. Under rational expectations, the price of an asset being fully equal to its discounted expected fundamental value (present value theory), we obtain the following relation:

$$P_t = \frac{1}{1+R} E_t(P_{t+1} + \psi_{t+1}), \tag{4.1}$$

where R is the constant discount rate (R > 0), P_t is the observed asset price at time t, and $\psi_t$is its discounted expected fundamental value, which takes the form of the real dividend/convenience yield received for owning the asset between t-1 and t.

When $t+n$ is far in the future, $\left[\frac{1}{(1+R)^n} E_t(P_{t+n})\right]$ does not affect P_t as:

$$\lim_{t+n\to\infty} \frac{1}{1+R} E_t(P_{t+n}) = 0. \tag{4.2}$$

This implies that n periods forward, Equation (4.1) can be rewritten as follows:

$$P_t^* \equiv E_t\left[\sum_{i=1}^{n} \frac{1}{1+R} (\psi_{t+n})\right] \tag{4.3}$$

The right-hand side of Equation (3) is also called the market fundamental solution. But if the terminal condition (3) does not hold, which means that a gap exists between the market fundamental solution and the actual price, an additional "bubble component", B_t, has to be added to the solution to equation (4.1):

$$P_t = P_t^* + B_t \,. \tag{4.4}$$

In that case, the so-called "market fundamental solution", P_t^*, is only the "fundamental component" of the price. And B_t is any random variable that satisfies the condition:

$$B_t = \frac{1}{1+R} E_t(B_{t+n}). \tag{4.5}$$

The bubble component is hence included in the price process and is expected to be present in the next period with an expected value times (*1* + *R*) its current value. Being then fully in line with the rational expectation framework, the bubble component maybe termed a “rational bubble”.

A series of indirect tests derives from testing the validity of the market fundamental solution. The variance bound test proposed by Shiller (1981), where the variance of the observed asset price should exceed the bound imposed by the variance of the fundamental value in the case of a rational bubble, is one of the earliest methods ever created. Though, it was not originally designed to test for the presence of bubbles, it can be used that way. A limitation is that emphasis is put on volatility, an element which can also be caused by variations in expected returns or by investment “fashion”, as pointed out by West (1987, 1988).

Campbell and Shiller (1987) leverage on the bubble literature and introduce another indirect method of bubble detection based on unit root testing. They anchor their approach on the idea that if a gap exists between the asset price and its fundamental value, this gap will exhibit an explosive behavior during the process of bubble formation. They identify two scenarios that can suggest strongly the presence of a rational bubble: (i) when the asset price is non-stationary in level but the fundamental value is, or (ii) when both the asset price and the fundamental value are non-stationary. In that second case, however, a co-integration test is needed; if the asset price and its fundamental value are co-integrated, and hence have co-movement in the long run, their non-stationary behavior is not a sign of the presence of a bubble. Diba and Grossman (1988) demonstrate that the idea of Campbell and Shiller (1987), i.e. the identification of explosive behavior in the gap between the asset price and its fundamental value, is sufficient to prove the existence of bubbles.

In spite of their limitations, pointed out by Evans (1991), left-tailed unit-roots tests and co-integration tests have been the go-to approach in bubble detection. Evans (1991) finds that that these tests fail to detect explosive bubbles when there are periodically collapsing bubbles in the time series because they are unable to differentiate between a periodically collapsing bubble trend and a stationary process. That is, collapsing bubbles show in the data “break” the non-stationary characteristics of the sample.

Therefore, a time series that contains several bubbles can be interpreted by the standard left-tailed unit-root test as a stationary series which points to the wrong conclusion that the data contains no bubble.
The limitations of the left-tailed unit root tests[19] have been taken into account in a recent and direct bubble testing approach developed by Phillips et al. (2011).[20] The authors arrange a right-tailed ADF test instead of the standard left-tailed test. While the left-tailed ADF test and the right-tailed ADF test used by Philips et al. (2011) both test the null hypothesis of a unit root behavior, their alternative hypotheses diverge: "stationary behavior" for the former and "mildly explosive" for the latter. Therefore, by looking in the data for direct evidence of non-linear explosive behavior Phillips et al. (2011) avoid the risk of misinterpreting a rejection of the null hypothesis as due to stationary behavior and hence they overcome the problem of recognizing a periodically collapsing bubble.

In the following we will apply the approach developed by Phillips et al. (2011) in identifying price bubbles to test whether the recent increase in the price of artwork can be characterized as a bubble (i.e., non-linear explosive behavior in our dataset). As Equation (4.4) indicates, the price of the assets tested here (P_t) is made of two components: a fundamental component (P_t^*) and a bubble component (B_t) so that:

$$P_t = P_t^* + B_t, \text{ where } P_t^* \equiv E_t\left[\sum_{i=1}^{n} \frac{1}{1+R}(\psi_{t+n})\right] \text{ and } B_t = \frac{1}{1+R} E_t(B_{t+n}).$$

The statistical properties of P_t^* and B_t determine those of P_t so that if ψ_t is an integrated process of order 1 - i.e. I(1) - P_t^* is also an I(1) process. If no bubble exists, B_t= 0. And the properties of P_t are only determined by those of P_t^*. However, if $B_t \neq 0$, current prices will exhibit explosive behavior, as B_t

[19] Left-tailed unit-root tests are generally standard Augmented Dickey-Fuller (ADF) tests (see Dickey and Fuller (1979)).
[20] Though first empirically implemented by Phillips et al. (2011) for bubble detection in stock prices, their methodology has also been used in the real estate literature (Yiu et al. (2012)), food commodities (Etienne et al. (2013)) and other agricultural commodities (Areal et al. (2013)).

reflects a stochastic process in which the expected value of next period's value is greater than or equal to the current period's value (Areal et al. (2013)).

Therefore, we implement a right-tailed version of the standard ADF test for a unit root (which means B_t= 0) against the alternative of an explosive root ($B_t \neq 0$, right tail):

$$H_0 : \delta = 1$$

$$H_1 : \delta > 1$$

where δ is the estimated first order regression coefficient from the following equation:

$$y_t = \mu + \delta y_{t-1} + \sum_{j=1}^{J} \phi_{r_w}^{j} \Delta y_{t-1} + \varepsilon_t \tag{4.6}$$

with μ being an intercept, J the lag order, r_w the sample window size and ε_t the error term.

Instead of estimating the model (4.6) as a whole (as in a standard ADF test) or by using rolling windows of a fixed, pre-determined size, we use forward recursive calculations of the ADF statistics with an expanding window, which not only provides more accuracy in case of multiple bubbles but also allows us to date-stamp the origination and collapse of the bubble. This method based on forward recursive regressions is called "sup ADF test" (referred to hereafter as "SADF test").

In forward recursive regressions, model (4.6) is estimated repeatedly by least squares, using subsets of the sample data incremented by one observation at a time. The sample size is normalized to 1 (R = 1), which yields a sample interval of $[0,1]$. The first observation in our sample is set as the starting point of the estimation window: $r_1 = 0$; and the end point of our initial estimation window, r_2, is set according to our choice of minimal window size such that the initial window size is $r_w = r_2$. Therefore, the first regression produces an ADF statistic denoted ADF_{r_1}. The model being estimated while incrementing the window size by one observation at each pass, the last estimation is based on the whole sample (i.e. r_w= 1) and the corresponding statistic is ADF_1.

Even though, the SADF is very similar in nature to an ADF test, the critical values for testing the null hypothesis differ since the right tail of the distribution is needed. Following the latest update of Philips et al. (2013), we derive the critical values from Monte-Carlo simulations, using the following random walk process with an asymptotically negligible drift:

$$y_t = dT^{-\eta} + \theta y_{t-1} + e_t, e_t \sim N(0,1), \tag{4.7}$$

where d, η, and θ are constant, T is the sample size, and e_t is the error term.

We set the significance level of the critical values at 90%, 95% and 99% and obtain the corresponding *t*-statistics after 2,000 replications (see Table 4.2).

In order to date-stamp the origination and termination of a bubble - in the case where the null hypothesis is rejected, we match the time series of the recursive *t*-statistic ADF_r against the right-tailed critical values of the asymptotic distribution of the standard ADF *t*-statistic. Let r_e be the origination date and r_f the collapse of some explosive behavior in the data, the estimates of these dates are given by:

$$\hat{r}_e = inf_{r \geq r_2} \left\{ r : ADF_r > cv_{\beta_n}^{adf}(r) \right\}, \tag{4.8}$$

$$\hat{r}_f = inf_{r \geq \hat{r}_e} \left\{ r : ADF_r < cv_{\beta_n}^{adf}(r) \right\}, \tag{4.9}$$

where $cv_{\beta_n}^{adf}(s)$ is the right side critical value of ADF_s that corresponds to a significant level of β_n, which we set at the 1% level. When plotted on a graph, the estimated starting point of a bubble is the first chronological observation for which ADF_s crosses upward the corresponding critical value. And conversely, the estimated termination point of a bubble is the first chronological observation for which ADF_s crosses downward the corresponding critical value.

4.3 Data and Art Price Index Construction

We construct a sample of repeat sales from the Blouin Art Sales Index (BASI), an online database that provides data on artwork sold at auction by over 350 auction houses worldwide.[21] The BASI database is presently the largest known database of artwork, containing roughly 4.6 million pieces by more than 225,000 individual artists for the period 1922 to 2013. We focus solely on paintings, which represent 2.7 million works of art in the database.

For each auction record, the database contains information on the artist, the artwork and the sale. We observe the artist's name, nationality, year of birth and year of death (if applicable). For the artwork, we know its title, year of creation, medium, size and style, as well as whether it is signed or stamped. For the sale, we have data on the auction house, date of the auction, lot number, hammer price (the price for which the artwork was sold, excluding any premiums paid from the buyer to the auction house, converted to U.S. dollars at the prevailing spot price) and whether the artwork has been "bought in" or was withdrawn.[22]

It goes without saying that art markets differ substantially from financial markets, and this potentially limits the strict applicability of well known financial techniques. Investing in art typically requires extensive knowledge of art and the art market, and substantial capital to acquire the work of well-known artists. The market is highly segmented and dominated by a few large auction houses, and only a small number of works are presented for sale throughout the year. Art comes with risk, deriving from both the physical risks of fire and theft and the possibility of reattribution to a different artist, so the cost of

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

insurance can be high. While auction prices represent the value of art work, they also represent a complex and subjective set of beliefs based on past, present and future prices, individual tastes and fashion.

Paintings are heterogeneous assets and a variety of physical and non-physical characteristics cause a painting to be unique. To construct our four individual art indices, Impressionist and Modern, Post-war and Contemporary, American, and Latin American, we follow standard hedonic modeling to separate these characteristics that determine the price of a painting. The dependent variable in our hedonic model is the natural logarithm of the sale price in USD. The independent variables used in our model describe the following characteristics: medium, auction house, surface, signature, estimate price, living status, artist reputation and sale date. A major disadvantage that comes with the regression of the hedonic pricing model is multicollinearity. A high correlation between the considered variables increases the number of standard errors of the regression coefficient. In order to overcome the problem of multicollinearity, each dummy variable contains a reference variable, which is deleted from the sample.

Sales date: These dummy variables are based on the sales date of the paintings. Each dummy variable represents one year from 1970 to 2013. A value of 1 indicates the painting is created in period *t* using the respective art index.

Auction house: According to de la Barre et al. (1994) the more renowned auction houses have a positive significant effect on the price of an individual painting. They reason that the more established and famous auction houses will offer the "best" work, while the less familiar and smaller auction houses will have lower quality paintings (see also, Renneboog and Spaenjers (2013)). We assign those auction houses a separate dummy variable, i.e. 1 indicates that the painting was auctioned by one of these leading auction houses: Christie's London or New York and Sotheby's London or New York. We can expect that the coefficient estimates will have a positive sign, assuming that the better and more expensive artists (paintings) will be auctioned by those auction houses.

Medium: We know from the seminal literature on the art market that oil on canvas is the most used material by painters and has the highest sales prices. Therefore we will specify oil on canvas as the reference variable in our hedonic regression model; we assign other media a separate dummy variable. The dummy D_{it} will have a value of 1, when one of the dummies has the appointed medium. We expect negative coefficient estimates since the reference variable oil on canvas is assumed to fetch the highest prices.

Surface: The variable surface explains the impact of the size of a painting and is calculated as the width multiplied by the height of the painting. These continuous surface values are logged in the hedonic regression model.

Signature: Anderson (1974) explains that the strength of the attribution towards the painter is a significant feature of the sales price. Paintings who are signed by the painter are more expensive than unsigned pieces. A dummy value of 1 indicates that the artist did not sign the painting. We expect that signed paintings are more valuable than unsigned painters and will thereby have positive coefficients.

Living status: Production of paintings will halt when the respective artist dies. Since the artist is no longer able to create artwork, one might assume that the value of the considered artist's paintings will increase. However, if the artist is not famous by the time of their death, there is also a chance that the value of her paintings will decrease given that the artist is no longer able to promote her work. We specify the dummy variable as follows: a value of 1 indicates that an artist is alive. Due to these contrarian explanations regarding the living status of an artist, it is expected that the coefficient will not be highly significant or significant at all.

To make use of our time dummy variables and to perform an OLS regression on the pooled data from all available sales, we construct the following hedonic regression model:

$$\ln P_{it} = \alpha + \sum_{j=1}^{z} \beta_j X_{ij} + \sum_{t=0}^{\tau} \gamma_t D_{it} + \varepsilon_{it} \qquad \varepsilon \sim N(0, \sigma^2), \tag{4.7}$$

where P_{it} represents the price of painting *i* at time *t*, α is the regression intercept, β_j is the coefficient value of quality characteristic *x*, *Xij* is the quality characteristic value of the painting, the antilog of λ_t reflects the coefficient value for the time dummy, and D_{it} represents the time dummy variable, which has the value of 1 when the paintings was sold in the considered time period *t.*

The estimated coefficients on the time dummies, i.e., the outcome from the hedonic regression model, are used to create the four different art price indices over the period 1970 to 2013. The Impressionist and Modern, Post-war and Contemporary, American, and Latin American art prices indices are computed using the following equation:

$$Index_{t+1} = \frac{\mathrm{Exp}\ (\mathrm{Y}_{t+1})}{\mathrm{Exp}(\mathrm{Y}_t)} \tag{4.8}$$

The antilog (or exponential) of the sequence of time dummies is taken. We set the first year, i.e. the base year to 100 and compute the relative changes to this base year for the subsequent years. Hence, the art price indices have been calculated and are presented in Figure 4.1 and summary statistics in Table 4.1.

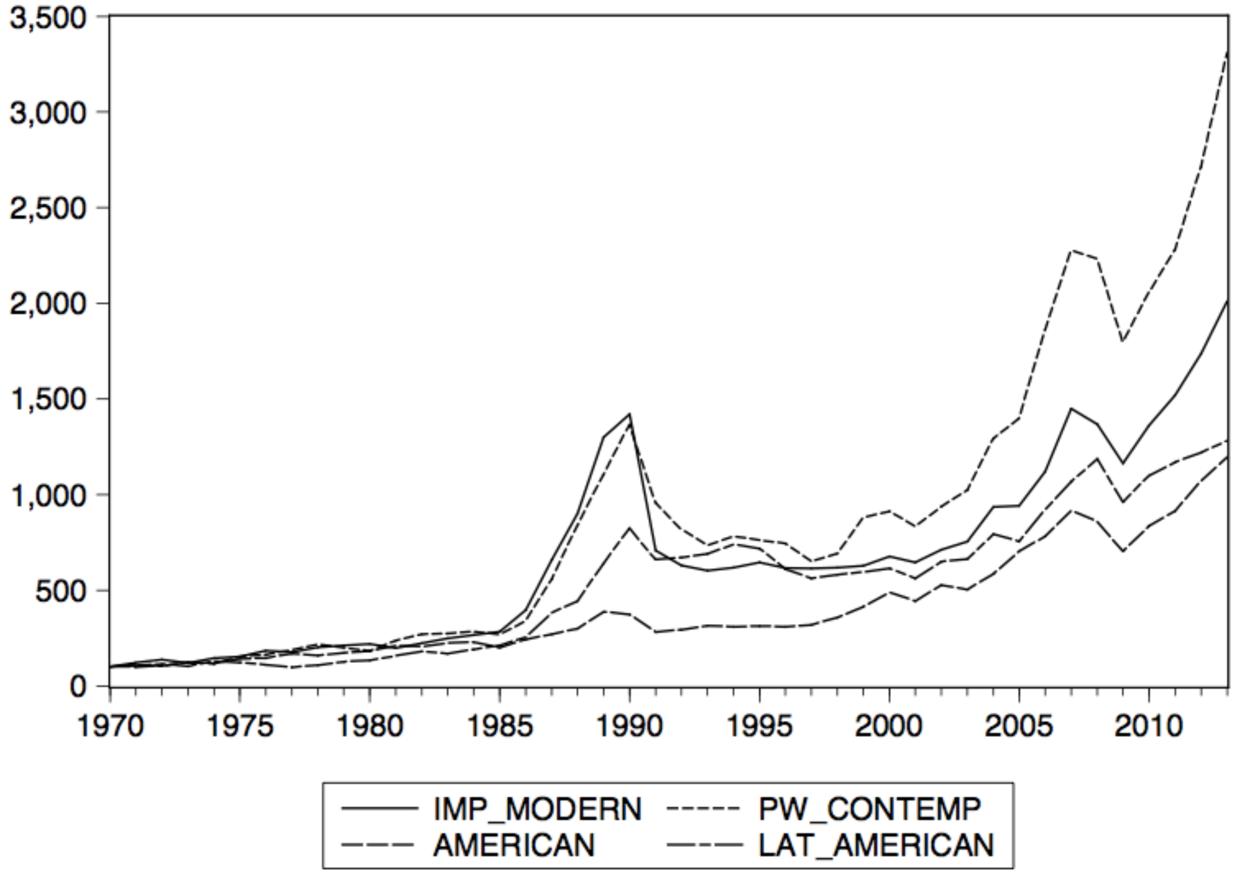

Figure 4.1: Art Market Indices. This figure shows the price indices for Impressionist & Modern, Post-War & Contemporary, American, and Latin American markets from 1970 to 2013. All series are normalized to 100 in 1970 (annual data).

Table 4.1: Summary Statistics.

This table provides the summary statistics of the four distinct art market indices in the Blouin Art Sales Index (BASI) data set over our sample period from 1970 to 2013. Based on 2.7 million works of art, the indices are constructed by using standard hedonic modeling with the following variables: medium, auction house, surface, signature, estimate price, living status, artist reputation and sale date. All annual series are normalized to 100 in 1970.

Variables	Observations	Mean	Standard Deviation	Min	Max
Impressionist & Modern	44	675.82	499.25	100	2,011
Post-War & Contemporary	44	889.14	796.47	100	3,311
American	44	399.45	284.50	99	1,197
Latin American	44	533.20	369.55	96	1,282

Figure 4.1 and Table 4.1 show that all four indices have risen dramatically over the last four decades, reaching levels in 2013 that range from 1,200 to 3,300 points. We can see that the Post-war and Contemporary market has been the most profitable segment but that each followed a similar trend with two identifiable jumps between 1985 and 1990 and between 2005 and 2009.

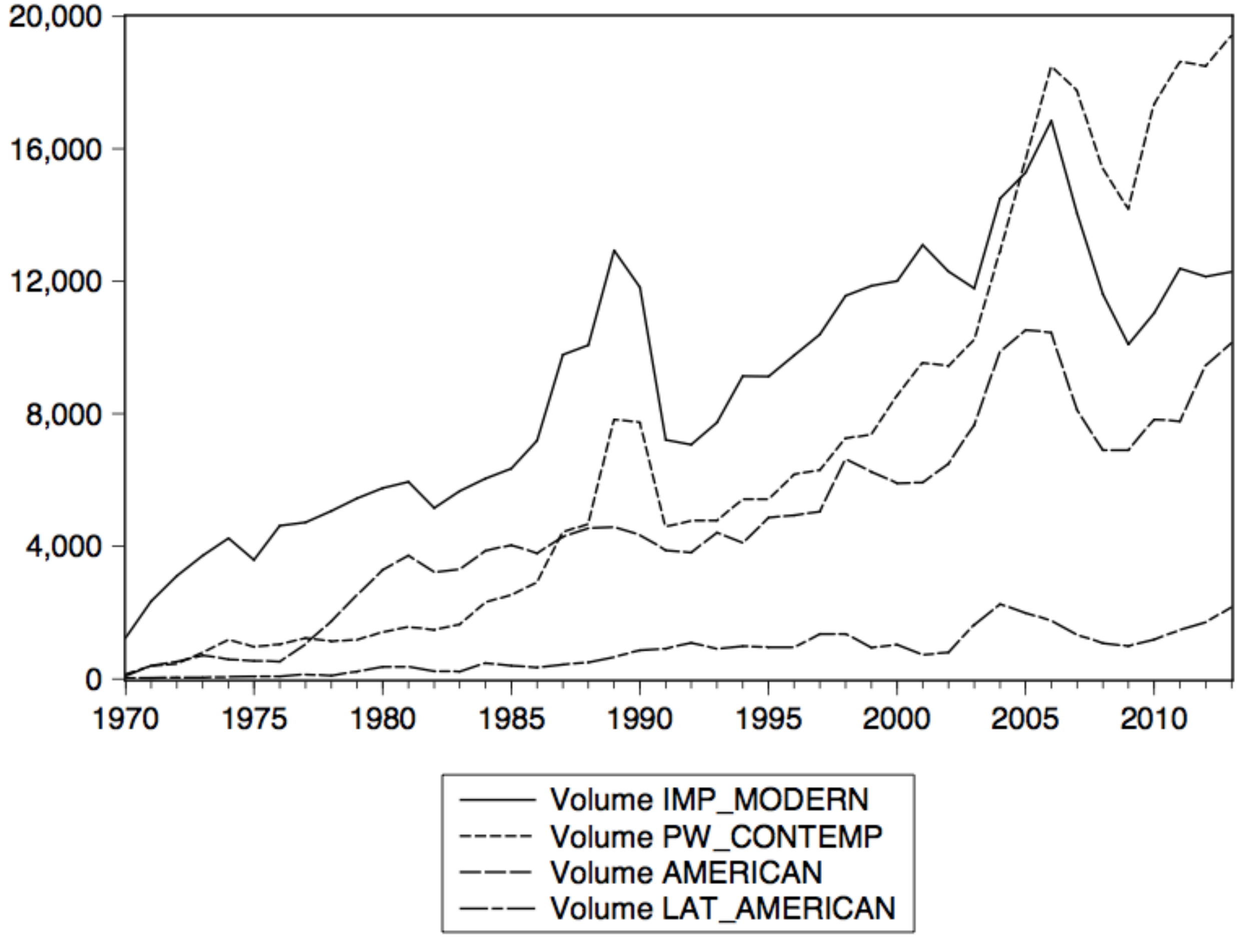

Figure 4.2: Art Market Volumes. This figure shows the total number of sales for each price index: Impressionist & Modern, Post-War & Contemporary, American, and Latin American markets from 1970 to 2013 (annual data).

Figure 4.2 displays the volumes (number of trades) over the period 1970 to 2013 for the distinct four art price indices and shows that the number of trades has increased by a factor of more than 80. A closer look also indicates that the "popularity" of segments has changed over time: while the Impressionist and Modern market had dominated the art market as a whole until the early 2000s, it is now the second segment in terms of volume, behind the Post-war and Contemporary market, which had already passed the American segment in the late 1980s.

4.4 Bubble Detection

We ran the test described in Section 4.2 for each of our four art market segments (Impressionist and Modern, Post-war and Contemporary, American, and Latin American). The summary output of the SADF test is displayed in Table 4.2 and shows that the critical values always appear below the *t*-stats of every segment. Therefore, we reject the null hypothesis of a unit root for each of them.

Table 4.2: Test for explosive behavior in the art market from 1970 to 2013.
This table reports the SADF tests of the null hypothesis of a unit root against the alternative of an explosive root, where the initial time-window is set to 10 observations and the significant level set to 99 percent. The series are the Impressionist & Modern art market, the Post-War & Contemporary art market, the American art market, and the Latin American art market. The sample period is 1973 to 2013 with 44 annual observations. The critical values for the SADF test are obtained by Monte-Carlo simulation with 2,000 replications.

	Time Series	**t-Statistics**
	Impressionist & Modern	9.818
	Post-War & Contemporary	6.371
	American	3.176
	Latin American	7.942
Test critical values for the explosive alternative:	99% level	1.959
	95% level	1.268
	90% level	0.983

Figures 4.3, 4.4, 4.5 and 4.6 present the date stamp procedure for the SADF test. They include the time series of the index, the ADF statistics sequence and its corresponding 1% critical values. In each case, the graphs confirm the results found in Table II since there is always at least one point in time between 1970

and 2013 where the ADF sequence juts out over its corresponding critical value, which clearly shows an explosive root and, therefore, the existence of at least one bubble in every subsample.

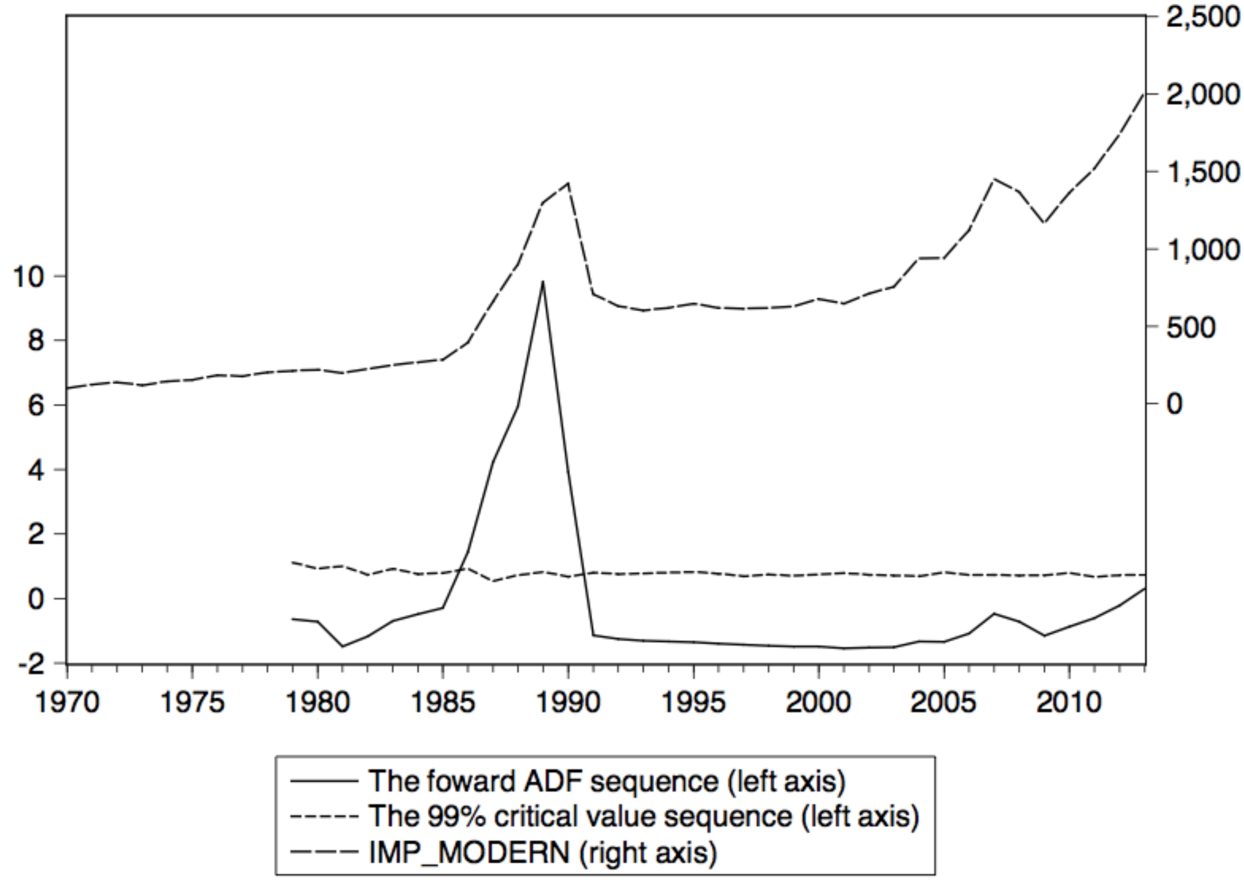

Figure 4.3: SADF Test and Date-Stamp, Impressionist & Modern Art Market. This figure represents graphically the SADF tests of the null hypothesis of a unit root against the alternative of an explosive and shows the time series of the SADF t-Statistic sequence for the Impressionist & Modern index on the left-hand (solid line), the corresponding sequence of critical values at the 99% level on the left-hand (small striped line) and the price index on the right-hand (large striped line) from 1970 to 2013. SADF t-Statistics are obtained from forward recursive calculations with an expanding window (initial size: 10 observations).

More specifically, the results presented in Figure 4.3 indicate (at the 1% critical level) the presence of only one bubble in the Impressionist and Modern art price index, which we date back to 1986 (using the date-stamping procedure explained in Section 2) until it burst in 1991, with a highest point reached in 1989. This finding is in line with the historical reality of the market, particularly as the growth and collapse of the 1986-1991 economic bubble in Japan which had fueled the art market with overconfidence and speculation. Japanese investors invested massively in international art markets in the late 1980s, funded by a "wealth effect" that gave them access to loans backed by the unsustainable collateral value of land and real estate prices in an overheated economy. When the uncontrolled credit expansion ended and the economy was no longer sustained by an excessively loose monetary policy, prices of real estate and land started to decrease. This obliged many Japanese investors to sell their holdings in art, often at a considerable discount (Hiraki et al (2009)).

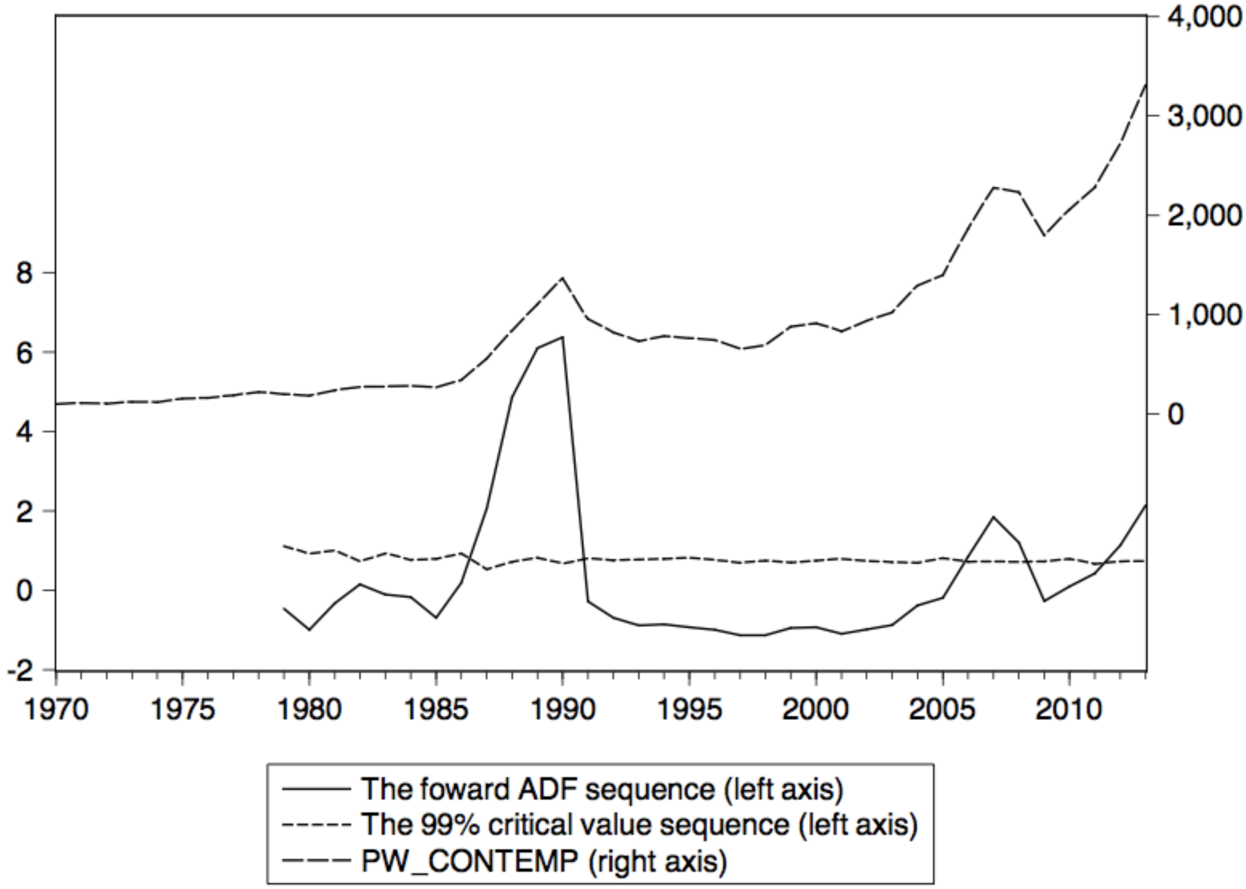

Figure 4.4: SADF Test and Date-Stamp, Post War & Contemporary Art Market. This figure represents graphically the SADF tests of the null hypothesis of a unit root against the alternative of an explosive and shows the time series of the SADF t-Statistic sequence for the Post-war and Contemporary index on the left-hand (solid line), the its corresponding sequence of critical values at the 99% level on the left-hand (small striped line) and the price index on the right-hand (large striped line) from 1970 to 2013. SADF t-Statistics are obtained from forward recursive calculations with an expanding window (initial size: 10 observations).

When looking at Figure 4.4, the Post-war and Contemporary subsample and its corresponding art price index, the same bubble is identified (here between 1986 and 1991 with a peak in 1990) and two others also appear to be present in the data: between 2006 and 2008 with a maximum reached in 2007 (which corresponds to the pre-financial crisis period) and from 2012 onwards (the recursions of 2009, 2010 and 2011 give upward moving *t*-stats that approach the critical values but the *t*-stats only exceed the critical values in 2012).

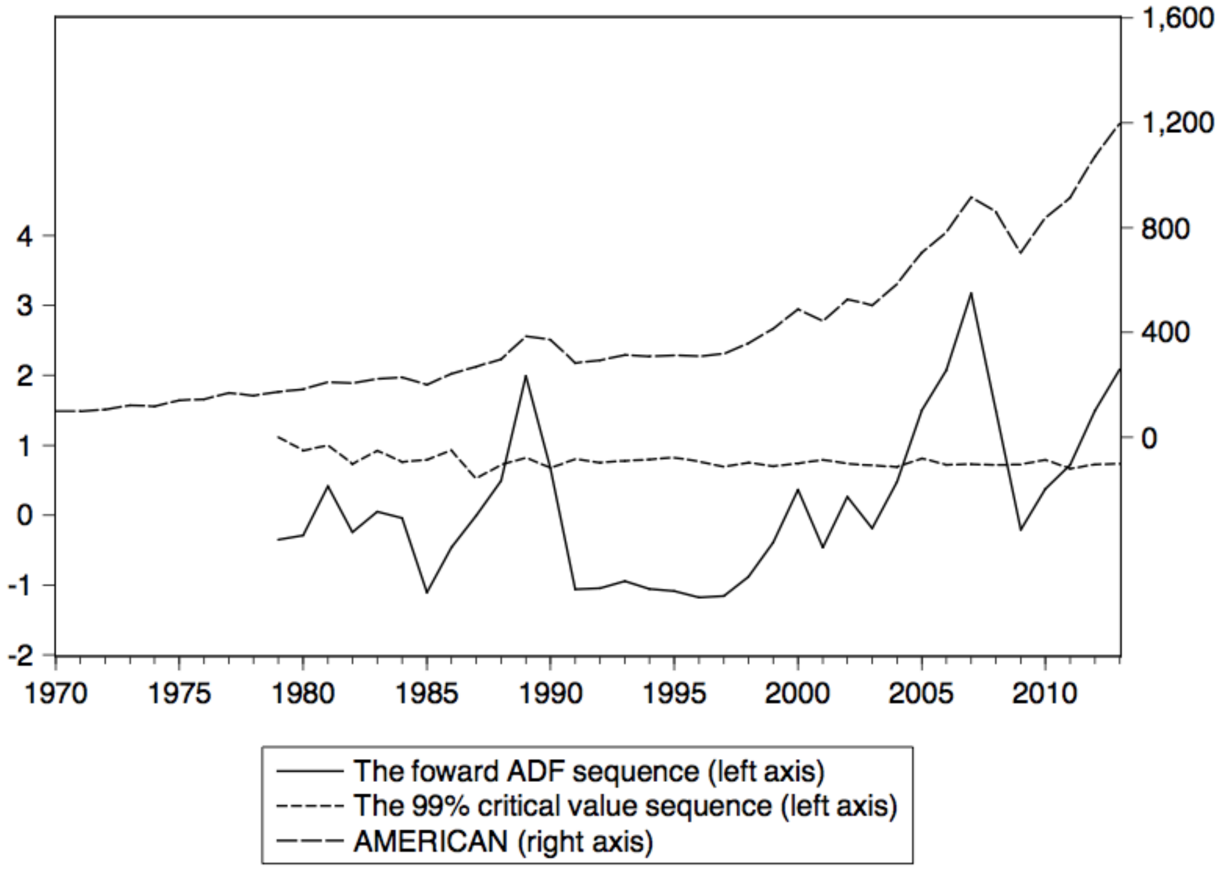

Figure 4.5: SADF Test and Date-Stamp, American Art Market. This figure represents graphically the SADF tests of the null hypothesis of a unit root against the alternative of an explosive behavior and shows the time series of the SADF t-Statistic sequence for the American index on the left axis (solid line), the its corresponding sequence of critical values at the 99% level on the left axis (small striped line) and the price index on the right axis (large striped line) from 1970 to 2013. SADF t-Statistics are obtained from forward recursive calculations with an expanding window (initial size: 10 observations).

For the American art price index, shown in Figure 4.5, the test detects similar explosive behavior in recent years (from 2011 onwards) to the one observed in the Post-war and Contemporary market segment. The previous 2006-2008 bubble in the Post-war and Contemporary market seems to start even earlier in the American market (2005-2008 with a maximum in 2007) but the 1990 bubble appears much shorter in length (1988-1991, peak in 1989) than for both the Impressionist and Modern market and Post-war and Contemporary market (both 1986 -1991).

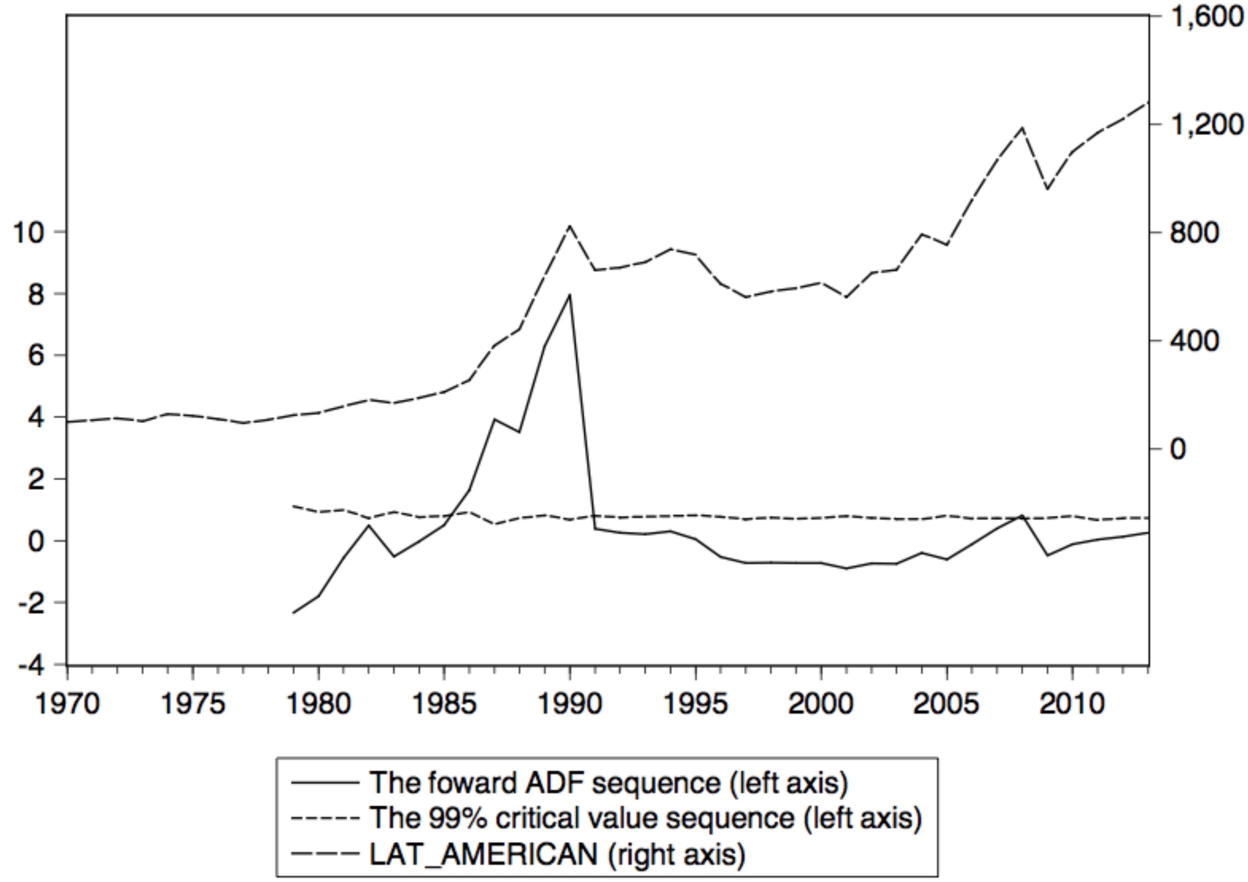

Figure 4.6: SADF Test and Date-Stamp, Latin American Art Market. This figure represents graphically the SADF tests of the null hypothesis of a unit root against the alternative of an explosive behavior and shows the time series of the SADF t-Statistic sequence for the Latin American index on the left-hand (solid line), the its corresponding sequence of critical values at the 99% level on the left-hand (small striped line) and the price index on the right-hand (large striped line) from 1970 to 2013. SADF t-Statistics are obtained from forward recursive calculations with an expanding window (initial size: 10 observations).

Figure 4.6 plots the test on the Latin American data and only shows the presence of the 1990 bubble (1986-1991, reaching a high point in 1990).

In other words, all market segments exhibit the same 1990 bubble but only the “Post-war and Contemporary” and the “American” market segments seem to have been affected by the 2008 bubble and the recent one, which started at the end of 2011. This 2011 one is according to our findings still in the mania phase of its formation (see the definition by Kindleberger and Aliber, 2005), i.e., it has not yet reached its peak and might very likely continue for another couple of years.

We observe that even though there exist some discrepancies amongst the origination dates of bubbles in the four studied subsamples, the termination dates coincide always perfectly. All bubble termination dates in all art market segments are in exact lockstep, which eliminates the possibility of arbitrages amongst different segments. Moreover, the particularities of artworks prevent investors from “riding” a bubble by rapidly trading a particular artwork, because such leads to a noticeable reduction in value (see Kräussl, 2013).

4.5 Conclusion

Our study examines the question of the existence of explosive behaviors in four different segments of the art market: Impressionist and Modern, Post-war and Contemporary, American, and Latin American.

The empirical findings suggest there is strong evidence of a current speculative bubble. Moreover, the mania phase appears to be on-going (after beginning in late 2011) both in the Post-war and Contemporary and in the American art market segments. We found that these two markets had already experienced speculative activity around 2006 - 2008. The absence of such phenomena in the other markets in our study suggests the Contemporary and the American are most likely to develop bubble-type behaviors.

The discovery of the well-known 1990 bubble which was present in every market, shows the degree of accuracy of the methodology developed by Phillips et al. (2011) and therefore gives even more robustness to the results stated above. Future research needs to concentrate on higher-frequency time series and more markets, as the data becomes available.

As the specificities of the art market make it almost impossible to “ride a bubble”, the analysis presented here finds its main application in the sell-side of the market. For instance, our results can be used by professionals (auction houses, art dealers, collectors) to better target the date of a sale or auction.

Chapter Five

General Conclusion

In the economic world, the perfect efficiency of markets has long been considered a truism. Yet the idea of the predictable nature of markets remains appealing to many investors. In this thesis, we examine three: the stock market (Chapter II), the options market (Chapter II and III) and the art market (Chapter IV).

The fact that irrationality and noise exist in the stock market has been well documented since the seminal work of Kahneman and Tversky (1974). On the other hand, the options market, where prices are linked strongly to a given formula (Black-Scholes), would appear to be the ideal candidate for having the most predictable prices and volatility levels. But even in this very special market, arbitrageurs have difficulties detecting opportunities and driving prices back to their intrinsic values. Our main conclusions are that the noise from the stock market is transmitted into the index options market, therefore spreading the irrationally of the stock market into the options market. We find that when noise traders are bearish (bullish) on a particular day, leading to flows out of (into) US equity funds, volatilities of S&P500 index options tend to increase (decrease) significantly on the following day. In line with other evidence, the effects are typically more pronounced for out-of-the-money options and short-term options. Furthermore, our findings suggest that the shape of the option smirk is caused partly by noise traders active in the equity market.

However, we also show that irrationality is not necessary where it is assumed to be. Stein (1989) demonstrates that long-term options have a tendency to react more to change of implied volatility than short-term options and do not display the smoothing which would be expected by rational behavior. Given the observed strong mean-reversion in volatility, Stein (1989) interpreted the results as evidence for overreaction in the options market, labeling it an “anomaly”. But we predict theoretically in this thesis that this stylized fact can be explained rationally by investors’ risk aversion. Long-term volatility should react more strongly to changes in short-term volatility in periods when investors are highly risk averse,and risk-neutral volatility is highly persistent. In contrast, in periods of low risk aversion, long-term volatility should react less strongly to changes in short-term volatility, because risk-neutral volatility is less persistent. We verify these theoretical predictions empirically and find the expected results: in periods of high risk aversion, long-term volatility reacts strongly to changes in short-term volatility, which can be explained by the high persistence of risk-neutral volatility in that period. The effect cannot be observed in periods of low risk aversion, due to a less persistent volatility process.

Finally, we suggest strongly that the well-known behavior biases that are displayed by actors in financial markets are also likely to be found in a non-financial setting. Our study of the art market shows that it is possible to see such a market as being like any other humanly-operated market. This leads us to conclude that there is strong evidence of both the Post-War and Contemporary and the American art market segments currently experiencing a speculative bubble, which began forming in late 2011 and is now in the mania phase of its formation. This finding could help make decisions in an asset allocation framework where art is seen as an alternative asset.

Overall, our results have implications for risk management, derivative pricing, and asset allocation.

Further research in the field of behavioral finance needs to identify the reasons why such a transmission of behavioral biases from the stock market to the options market is observed. The field of art and finance is also promising for further research, in particular investigating more deeply the microstructure of the art market and working to extend to more markets the bubble detection test. Another line of research would be to better understand and then attempt to forecast the end of such explosive behaviors.

References

Chapter One

Simon, A., 1997, Models of bounded rationality*,* empirically grounded economic reason, Vol 3, Cambridge: The MIT Press.

Rubinstein, M., (2000) Rational markets: yes or no? the affirmative case, *Financial Analysts Journal* 57, 15-29.

Chapter Two

Acharya, V., Johnson, T., 2007, Insider trading in credit derivatives, Journal of Financial Economics.

Baker, M., Wurgler, J., 2006, Sentiment and the cross-section of stock returns. Journal of Finance 61, 1645-1680.

Baker, M., Wurgler, J., 2007, Investor sentiment in the stock market. Journal of Economic Perspectives 21, 129-151.

Bakshi, G., Kapadia, N. Madan, D., 2003, Stock Return Characteristics, Skew Laws, and Differential Pricing of Individual Equity Options, *Review of Financial Studies* 16, 101-143.

Bams, D., Lehnert, T., Wolff, C., 2009. Loss Functions in Option Valuation: A Framework for Selection. Management Science 55 **(**5), 853-862.

Barberis, N., Shleifer, A., Vishny, R., 1998. A Model of Investor Sentiment. Journal of Financial Economics 49, 307-343.

Barone-Adesi, Giovanni, Engle, Robert F. and Mancini, Loriano, GARCH Option Pricing Model with Filtered Historical Simulation (January 2008). *Review of Financial Studies*, 2008.

Beaumont, R., van Daele, M., Frijns, B., Lehnert, T., and Muller, A. 2008. Investor sentiment, mutual fund flows and its impact on returns and volatility. Managerial Finance 34, 11, 772-785.

Ben-Rephael, Azi and Kandel, Shmuel and Wohl, Avi, 2011, Measuring Investor Sentiment with Mutual Fund Flows. Journal of Financial Economics, forthcoming.

Berndt, A., Ostrovnaya, A., (2008). Do equity markets favour credit market news over options market news? Carnegie Mellon University. Working Paper.

Black, F., Scholes, M., 1973. The Pricing of Options and Corporate Liabilities. Journal of Political Economy 81, 637-659.

Bollen, N., Whaley R., 2004, Does Net Buying Pressure Affect the Shape of Implied Volatility Functions?, The Journal of Finance, 59, 2, 711-753.

Brown, S.J., Goetzmann, W.N., Hiraki, T., Shiraishi, N., Watanabe, M, 2005. Investor Sentiment in Japanese and U.S. Daily Mutual Fund Flows. Working paper.

Chalmers, J.M.R., Edelen, R.M. and Kadlec, G.B. (2001), On the perils of financial intermediaries setting security prices: the mutual fund wild card option, Journal of Finance, Vol. 56, pp. 2209-36.

Chiu, H., Kini, O., 2011. Equity Issuances, Equity Mutual Fund Flows, and Noise Trader Sentiment. Working Paper. Georgia State University.

Chevalier, Judith, and Glenn Ellison, 1997, Risk taking by mutual funds as a response to incentives, Journal of Political Economy 105, 1167–1200.

Christoffersen, P., Heston, S., Jacobs, K., 2010. Option Anomalies and the Pricing Kernel. Working Paper.

De Long, J.B., Shleifer, A., Summers, L.H., Waldmann, R.H., 1990. Sentiment trader risk in Financial markets. Journal of Political Finance 98, 703-738.

Dumas, B., J. Fleming and R.E. Whaley (1998): ”Implied Volatility Functions: Empirical Tests”, *The Journal of Finance* 53, 2059-2106.

Edwards, F. and Zhang, X. 1998. Mutual funds and stock and bond market stability. Journal of Financial Services Research 13, 3, 257-282.

Frijns, B., Lehnert, T., Zwinkels R., 2010. Behavioral Heterogeneity in the Option Market. Journal of Economic Dynamics and Control 34, 2273-2287.

Frijns, B., Lehnert, T., Zwinkels R., 2012. Sentiment Trades and Option Markets, Working Paper, University of Luxembourg.

Garleanu, N., Pedersen, L.H., Poteshman, A.M., 2007. Demand-Based Option Pricing. NBER Working Paper 11843.

Goetzmann, W. N., Massa, M, Rouwenhorst. K. G, 2000. Behavioral Factors in Mutual Fund Flows. Yale School of Management Working Paper.

Greene, J.T. and Hodges, C.W. (2002), ‘‘The dilution impact of daily fund flows on open-end funds’’, Journal of Financial Economics, Vol. 65, pp. 131-58.

Han, Bing, Investor Sentiment and Option Prices. *Review of Financial Studies*, 2008.

ICI. 2011. ICI Fact Book, 51st ed. Investment Company Institute.

Ippolito, R. 1992. Consumer reaction to measures of poor quality: Evidence from the mutual fund industry. Journal of Law and Economics 35, 1, 45-70.

Lakonishok, Josef, Inmoo Lee, Neil Pearson, and Allen Poteshman, 2006, Option Market Activity, forthcoming Review of Financial Studies.

Lee, W. Y., Jiang, C. X., & Indro, D. C. (2002). Stock market volatility, excess returns, and the role of investor sentiment. Journal of Banking & Finance, 26(12), 2277-2299.

Mahani, Reza, and Allen Poteshman, 2004, Overreaction to Stock Market News and Misevaluation of Stock Prices by Unsophisticated Investors: Evidence from the Option Market, Working Paper, University of Illinois at Urbana-Champaign.

Newey, W., West, K., 1987. A simple positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica* 55, 703–708.

Poteshman, A.M., 2001. Underreaction, Overreactions, and Increasing Misreaction to Information in the Options Market. *Journal of Finance* 56, 851-876.

Poteshman, Allen, and Vitaly Serbin, 2003, Clearly Irrational Financial Market Behavior: Evidence from the Early Exercise of Exchange Traded Stock Options, *Journal of Finance* 58,37-70.

Sirri, E. and Tufano, P. 1998. Costly search and mutual fund flows. *Journal of Finance* 53, 5, 1589-1622.

Stein, Jeremy, 1989, Overreactions in the Options Market, *Journal of Finance* 44, 1011-1022.

Chapter Three

Ait-Sahalia, Y., Lo, A., 2000, Nonparametric Risk Management and Implied Risk Aversion. Journal of Econometrics, Vol. 94: 9-51.

Baker, M., Wurgler, J., 2006, Sentiment and the cross-section of stock returns. Journal of Finance, Vol. 61, 1645-1680.

Baker, M., Wurgler, J., 2007, Investor sentiment in the stock market. Journal of Economic Perspectives, Vol. 21, 129-151.

Bakshi, G., Kapadia, N., Madan, D., 2003, Stock return characteristics, skew laws, and differential pricing of individual equity options. Review of Financial Studies, Vol. 16, 101-143.

Bakshi, G., Madan, D., Panayotov, G., 2010, Deducing the Implications of Jump Models for the Structure of Crashes, Rallies, Jump Arrival rates and Extremes. Journal of Business and Economic Statistics. Vol. 28: 380-396.

Barone-Adesi, G., Engle, R. F., Mancini, L., 2008, A GARCH option pricing model with filtered historical simulation, Review of Financial Studies, Vol. 21, No. 3, 1223-1258.

Bekkour, L., Jin, X., Rasmouki, F., Wolff, C., 2013, Euro at Risk: The Impact of Member Countries' Credit Risk on the Stability of the Common Currency. University of Luxembourg, working paper.

Bollen, N., Whaley R., 2004, Does net buying pressure affect the shape of implied volatility functions? The Journal of Finance, Vol. 59, 2, 711-753.

Christoffersen, P., Heston, S., Jacobs, K., 2013, Capturing Option Anomalies with a Variance-Dependent Pricing Kernel. The Review of Financial Studies, Vol. 26, 8, 1963-2006.

Dumas, B., Fleming, J., Whaley, R.E., 1998, Implied volatility functions: empirical tests. Journal of Finance, Vol. 53, 2059-2106.

Han, B., 2008, Investor Sentiment and Option Prices. Review of Financial Studies, Vol. 21, 387-414.

Heston, S., 1993, A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options. The Review of Financial Studies, Vol. 6, number 2, 327-343.

Jackwerth, J.C., 2000, Recovering Risk Aversion from Options Prices and Realized Returns. The Review of Financial Studies, Vol. 13: 433-451.

Lehnert, T., Lin, Y., Wolff, C., 2013, Skewness Risk Premium: Theory and Empirical Evidence. CEPR Discussion Paper.

Poteshman, A. M., 2001, Underreaction, overreaction, and increasing misreaction to information in the options market. Journal of Finance, Vol. 56, No. 3, 851-876.

Shive, S., Shumway, T., 2009. Is the pricing kernel monotonic? University of Notre Dame. Working Paper.

Stein, J., 1989, Overreactions in the options market. Journal of Finance, Vol. 44, No. 4, 1011-1023.

Chapter Four

Areal, F., K. Balcombe, and G. Rapsomanikis, 2013. Testing for bubbles in agriculture commodity markets. University of Reading Working Paper, No. 48015.

Campbell, R., 2008. Art as a financial investment. *Journal of Alternative Investments* 10(4), 64-81.

Campbell, J., and R. Shiller, 1987. Co-integration and tests of present value models. *Journal of Political Economy* 95(5), 1062-1088.

Chanel, O., L. Gerard-Varet, and V. Ginsburgh, 1996. The relevance of hedonic price indices. *Journal of Cultural Economics* 20(1), 1-24.

De Aenlle, C., 2014, How investors can thrive in a lopsided art market, *New York Times*, April 29.

Diba, B., and H. Grossman, 1988. Explosive rational bubbles in stock prices. *American Economic Review* 78(3), 520-530.

Dickey, D., and W. Fuller, 1979, Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association* 74(366), 427-431.

Etienne, X., S. Irwin, and P. Garcia, 2013, Bubbles in food commodity markets: Four decades of evidence. *Journal of International Money and Finance* 42(C), 129-155.

Evans, G., 1991. Pitfalls in testing for explosive bubbles in asset prices. *American Economic Review* 81(4), 922-930.

Figuerola-Ferretti, I., and J. Gonzalo, 2010. Modeling and measuring price discovery in commodity markets. *Journal of Econometrics* 158(1), 95-107.

Frey B., and W. Pommerenhe, 1989. *Muses and markets: Explorations in the economics of the arts*. Basil Blackwell, Oxford.

Gérard-Varet, L., 1995. On pricing the priceless: Comments on the economics of the visual art market. *European Economic Review* 39(3-4), 509-518.

Hiraki, T., I. Akitoshi, D. A. Spieth, and N. Takezawa, 2009. How Did Japanese Investments Influence International Art Prices? *Journal of Financial and Quantitative Analysis* 44 (6), 1489-1514.

Kindleberger, C., and R. Aliber, 2005. *Manias, panics, and crashes: A history of financial crises*. John Wiley and Sons, Hoboken.

King, R., V. Smith, A. Williams, and M. van Boening, 1993. *The robustness of bubbles and crashes in experimental stock markets*. Oxford University Press.

Korteweg, A., R. Kräussl, and P. Verwijmeren, 2013. Does it Pay to Invest in Art? A Selection-corrected Returns Perspective. Stanford Graduate School of Business Working Paper.

Kräussl, R., 2013. To have and to hold? Databank, p. 104-5, *Art + Auction*, July-August 2013.

Kräussl, R., and C. Wiehenkamp, 2012. A call on art investments. *Review of Derivatives Research* 15(1), 1-23.

McAndrew, C., 2008. *The art economy: An investor's guide to the art market*. Liffey Press, Dublin.

Phillips, P., Y. Wu, and J. Yu, 2011. Explosive behavior in the 1990s NASDAQ: When did exuberance escalate asset values? *International Economic Review* 52(1), 201-226.

Phillips, P., S. Shi, and J. Yu, J., 2013a. Specification sensitivity in right-tailed unit root testing for explosive behaviour. Forthcoming in *Oxford Bulletin of Economics and Statistics*.

Phillips, P., S. Shi, and J. Yu, 2013b. Testing for multiple bubbles: Historical episodes of exuberance and collapse in the S&P 500. Cowles Foundation Discussion Paper 1914 .

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

Rooney, B., 2013. Is there a bubble in the art market? *CNN Money*, November 15.

Shiller, R., 1981. Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review* 71(June 1981), 421-436.

Stein, J., 1977. The monetary appreciation of paintings. *Journal of Political Economy* 85(5), 1021-1035.

West, K., 1987. A specification test of speculative bubbles. *Quarterly Journal of Economics* 102(3), 553-580.

West, K., 1988. Bubbles, fads and stock price volatility tests: A partial evaluation. *Journal of Finance* 43(3), 639-656.

Yiu, M., L. Jin, and J. Yu, 2012. Detecting bubbles in the Hong Kong residential property market: An explosive-patter approach. Hong Kong Institute for Monetary Research Paper No. 3

[i] This section is extracted from: Frijns, B., Irek, F., Lehnert, T., Martelin, N., 2012. Noise Trading and the Cross-Section of Index Option Prices. LSF Research Working Paper Series No.12-1

[ii] This section is extracted from: Lehnert, T., Lin, Y., Martelin, N., 2013. Stein Overreaction Puzzle: Option Anomaly or Perfectly Rational Behavior?. LSF Research Working Paper Series No.13-2

[iii] This section is extracted from: Kräussl, R., Lehnert, T., Martelin, N., 2014. Is There a Bubble in the Art Market?. Forthcoming in *The Journal of Wealth Management*