Stein's Overreaction Puzzle: Option Anomaly or Perfectly Rational Behavior?

THORSTEN LEHNERT*
Luxembourg School of Finance, University of Luxembourg

YUEHAO LIN
Luxembourg School of Finance University of Luxembourg

NICOLAS MARTELIN
Luxembourg School of Finance University of Luxembourg

November 2013

ABSTRACT

Empirical studies have shown that implied volatilities of long-term options react quite 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, those findings have been interpreted as evidence for overreaction in the options market. Focusing on a stochastic variance process in a rational expectation framework, we theoretically show that under normal market conditions the risk-neutral volatility dynamics are substantially more persistent than the physical one. As a result, the empirical observation of a strong reaction of long-term volatility would be consistent with perfectly rational behavior. We additionally show that the degree of persistence depends on investors' risk aversion. Using daily data on S&P 500 index options, we confirm previous findings for the 2000-2010 period, which is characterized by an overall moderate level of risk aversion. Once we identify periods of high and low risk aversion, in line with the predictions of our theoretical model, empirical results only hold for periods, when investors are highly risk averse. During periods of low risk aversion, results are insignificant. Robustness checks reveal that the results are remarkably stable over the complete term structure. Therefore, we provide strong evidence that the empirical observation is not overreaction, but in line with perfectly rational behavior.

Keywords: Option Markets, Overreaction, Rational Expectations, Mean Reversion, Volatility
JEL-Classifications: G08, G12

* Corresponding Author: Luxembourg School of Finance, University of Luxembourg, 4, rue Albert Borschette, 1246 Luxembourg, Luxembourg, tel +352466644-6941; fax +352466644-6835. E-mail address: thorsten.lehnert@uni.lu .

Stein's Overreaction Puzzle:
Option Anomaly or Perfectly Rational Behavior?

ABSTRACT

Empirical studies have shown that implied volatilities of long-term options react quite 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, those findings have been interpreted as evidence for overreaction in the options market. Focusing on a stochastic variance process in a rational expectation framework, we theoretically show that under normal market conditions the risk-neutral volatility dynamics are substantially more persistent than the physical one. As a result, the empirical observation of a strong reaction of long-term volatility would be consistent with perfectly rational behavior. We additionally show that the degree of persistence depends on investors' risk aversion. Using daily data on S&P 500 index options, we confirm previous findings for the 2000-2010 period, which is characterized by an overall moderate level of risk aversion. Once we identify periods of high and low risk aversion, in line with the predictions of our theoretical model, empirical results only hold for periods, when investors are highly risk averse. During periods of low risk aversion, results are insignificant. Robustness checks reveal that the results are remarkably stable over the complete term structure. Therefore, we provide strong evidence that the empirical observation is not overreaction, but in line with perfectly rational behavior.

I. INTRODUCTION

Proponents of the efficient markets hypothesis would claim that investors correctly incorporate new information into asset prices. Bayesian rationality is assumed to be a good description of investor behavior. Empirical studies are challenging this view. One interesting and robust stylized fact that emerges from the index options literature is the overreaction puzzle of Stein (1989), which was further investigated by Poteshman (2001) and more recently by Christoffersen et al. (2013).

Stein (1989) derives and empirically tests 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.

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 put options and from distant ATM call and 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 are challenging 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[1], 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[2]. Our study is building on their theoretical results.

Assuming a stochastic variance process in a rational expectation framework, we theoretically show 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 empirically verify 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 highly persistent risk-neutral volatility dynamics in that period. The effect cannot be observed in periods of low risk aversion, because of less persistent

[1] One-month maturity for short-term options and two-month maturity for long-term option

[2] In their model, the pricing kernel specification is monotonic in returns and also monotonic in volatility.

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 the paper is organized as follows. In the next section, we present the theoretical framework. In Section 3 and 4, we discuss the data and present the empirical analysis and results. Section 5 concludes.

II. THEORETCAL FRAMEWORK

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 that is 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's (1993) model

$$
\begin{aligned}
dS(t) &= \big(r + \mu v(t)\big)S(t)dt + \sqrt{v(t)}S(t)dZ_1(t) \\
dv(t) &= \kappa\big(\theta - v(t)\big)dt + \sigma\sqrt{v(t)}dZ_2(t)
\end{aligned}
$$

where $S(t)$ is the spot price; $v(t)$ is the instantaneous variance; r is the risk-free rate; μ relates to the equity risk premium; $Z_1(t)$ and $Z_2(t)$ are Wiener processes, where $Z_2(t)$ has a correlation ρ with $Z_1(t)$. The variance is assumed to revert towards a long run mean, θ, with a mean reverting speed κ. In a risk-neutral world, the variance follows the same mean reverting process but under risk-neutral measures

$$
\begin{aligned}
dS(t) &= rS(t)dt + \sqrt{v(t)}S(t)dZ_1^*(t) \\
dv(t) &= \kappa^*\big(\theta^* - v(t)\big)dt + \sigma\sqrt{v(t)}dZ_2^*(t)
\end{aligned}
$$

where $\kappa^* = \kappa + \lambda$ and $\theta^* = \kappa\theta/(\kappa + \lambda)$. Therefore, under the risk-neutral pricing probability, the variance reverts 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

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

Without loss of generality, we assume a two-period model. By maximizing utility over the two periods when the investor chooses to allocate his wealth on stocks or to consume, in equilibrium, we get a stock price process satisfying the martingale pricing condition. The pricing kernel, based on the consumption return state (c_T/c_t) is thus as follows

$$\frac{\pi_T}{\pi_t} = \beta \left(\frac{c_T}{c_t}\right)^{-\gamma}, with\ \gamma = \frac{-cU''(c)}{U'(c)}$$

where the coefficient β is a time preference function and γ 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).[3]

Depending on where Taylor expansion series is truncated of the stock price process, the pricing kernel takes 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

[3] 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).

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

where δ and η determine investor's return preference; ξ determines the variance preference; and γ is the constant relative risk aversion coefficient. Compared with the basic one, $\frac{\pi_T}{\pi_t} = \left(\frac{S_T}{S_t}\right)^{-\gamma}$, the more general pricing kernel takes into account not only risk aversion, but also investor's variance preference. Out of arbitrage conditions, 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 can be derived as (see Christoffersen et al. (2013))

$$\lambda = \rho\sigma\gamma - \sigma^2\xi$$

The correlation between stock market returns and its variance, ρ, is empirically found to be negative, $\rho < 0$. 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, $\lambda < 0$, being it from a pricing kernel based on kinds of varieties of stochastic models or from a 'by-product' of general equilibrium models.

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[4], investors require more compensation for risk, both terms, $\rho\sigma\gamma$ and $\sigma^2\xi$, 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 substantially lower rate of mean reversion, or a more persistent risk-neutral volatility process, as $\kappa^* = \kappa + \lambda$ and, therefore, $\kappa^* = \kappa + (\rho\sigma\gamma - \sigma^2\xi) < \kappa$.

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 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, the 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, 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

[4] 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 be increasing in the variance.

(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.

III. DATA & METHODOLOGY

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 successfully eliminates 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[5]:

$$IV_{i,j} = \alpha_0 + \alpha_1 delta_{i,j} + \alpha_2 delta_{i,j}^2 + \alpha_3 T_j \;\; + \alpha_4 T_j^2 + \alpha_5 delta_{i,j} \; T_j$$

where IV_{ij} denotes the Black-Scholes implied volatility and $delta_{i,j}$, the call delta of an option for the i-th exercise price and j-th maturity[6]. 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))[7]. In current 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 currency options). We derive model-free measures of the 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,

$$Var_t(T) = e^{rt} V_t(T) - \mu_t^2(T)$$

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

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

[5] 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.

[6] For put options, we use the corresponding call delta in the implied volatility regression.

[7] 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.

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

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

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 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 1 month to be short-term and 2 months to be long-term. Additionally, as a robustness check, we replicate the analysis using different pairs of “nearby” and “distant” options. Table I presents an overview of risk-neutral annualized volatilities for short- as well as long-term options.

[Table I]

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 both periods of low volatility levels (post-dotcom bubble years) and more volatile periods (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 max figures, in particular during periods of market stress, short-term volatilities can increase substantially more than long-term volatilities. Once we subdivide the sample period into periods of high and low sentiment[8], we find that the average volatilities are quite similar, but in high sentiment periods, volatilities can increase substantially, as can be seen from the max 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 is supporting our main hypothesis. In the next section, we proceed with the term-structure tests.

[8] 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.

IV. EMPIRICAL ANALYSIS

In order to test if investors are able to correctly incorporate new information 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[9]:

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

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

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

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 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[10].

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[11].

---

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

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

[11] 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.

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

Short-term options are assumed to be one-month options and long-term options are assumed two-months options. The difference between the two terns is assumed to be one month or 21 trading days. All regressions are standard OLS and the results are displayed in Table II.

[Table II]

Our results are in line with the common findings in the literature. The estimates displayed in Table II 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 would Stein (1989) considers to be “normal” reaction. However, interestingly, the coefficients presented in Table II, 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 1 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 2, which we would like to test empirically in the subsequent remaining part of the paper.

[Figure 1]

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[12]. 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[13]. 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.

[12] We have also constructed other sentiment indicators using the original Baker and Wurgler (2006; 2007) data, but the results qualitatively don't change.

[13] See Lehnert et al. (2013) for details.

[Table III]

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 dramatically weakened (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%. Generally, 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, substantially strengthens the argument in favor of a rational explanation.

As a robustness check, we run the same regressions over the complete term structure of risk-neutral volatility. In his analysis, Stein (1989) considered "nearby" options to have 1 month maturity and "distant" options to have 2 months maturity. However, it can be shown that the same term structure tests can be conducted for various combinations of "nearby" and "distant" options. In particular, we run the regressions for 2 months options versus 4 months options, for 3 months options versus 6 months options and for 4 months options versus 8 months. This examination also represents a robustness check of the Christoffersen et al. (2013) and Stein (1989) analysis. The results are presented in Table IV and V.

[Table IV-V]

Results suggest that all our previous findings hold, but get even stronger once we consider different parts of the term structure. All coefficients are significantly negative. 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 are remarkably stable over the complete term structure. Again, the two-regime regressions accommodate the data much better than the one-regime model, with R^2 increasing from 4% to 19%, from 5% to 25% and from 7% to 28%, respectively.

V. 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 are challenging this view. Building on a stochastic variance process in a rational expectation framework, we theoretically show 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 empirically verify 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.

REFERENCES

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., Lehnert, T., 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.

Table I
Options Volatilities - Summary Statistics

Summary statistics of options risk-neutral volatilities.

Maturity (months)	Mean	Stand. Dev.	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						
Maturity (months)	Mean	Stand. Dev.	Min	Max	AR 1	N
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						
Maturity (months)	Mean	Stand. Dev.	Min	Max	AR 1	N
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

Table II

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	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

Table III

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 (1 month). Vol_t^{2m} is the risk-neutral volatility of options with long-term maturity (2 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	
α_H	-0.016	0.005	-3.13	2560
β	-0.205	0.016	-12.98	
β_H	0.203	0.022	9.24	

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

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

Vol_t^{sm} is the risk-neutral volatility of options with short-term maturity s. Vol_t^{lm} is the risk-neutral volatility of options with long-term maturity l. Panel A shows the results assuming the short term to be 2 months (s=2) and the long term to be 4 months (l=4). Equivalently, Panel B and C present the results for 3 vs. 6 months, and 4 vs. 8 months, respectively. X corresponds to the time difference (in trading days) between long-term options and short-term options.

	Coefficient	Standard Error	t-Statistic	N
Panel A: 2 vs. 4 months				
α	0.012	0.004	3.50	2539
β	-0.155	0.0150	-1030	
Panel B: 3 vs. 6 months				
α	0.021	0.004	5.21	2518
β	-0.195	0.017	-11.79	
Panel C: 4 vs. 8 months				
α	0.035	0.004	8.34	2497
β	-0.235	0.017	-13.67	

Table V

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

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

Vol_t^{Sm} is the risk-neutral volatility of options with short-term maturity s. Vol_t^{lm} is the risk-neutral volatility of options with long-term maturity l. Panel A shows the results assuming the short term to be 2 months (s=2) and the long term to be 4 months (l=4). Equivalently, Panel B and C present the results for 3 vs. 6 months, and 4 vs. 8 months, respectively. 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: 2 vs. 4 months				
α	0.019	0.005	4.07	2539
α_H	-0.024	0.007	-3.67	
β	-0.308	0.019	-16.07	
β_H	0.333	0.028	12.04	
Panel B: 3 vs. 6 months				
α	0.023	0.005	4.67	2518
α_H	-0.030	0.007	-4.17	
β	-0.361	0.020	-18.31	
β_H	0.413	0.030	13.98	
Panel C: 4 vs. 8 months				
α	0.026	0.005	5.20	2497
α_H	-0.020	0.007	-2.63	
β	-0.364	0.020	-18.47	
β_H	0.387	0.031	12.58	

Figure 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 (1 month). Vol_t^{2m} is the risk-neutral volatility of options with long-term maturity (2 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 6 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.