Matthias Leiss, Heinrich H. Nax, Didier Sornette

Super-exponential growth expectations and the global financial crisis

Article (Accepted version) (Refereed)

Original citation:
Leiss, Matthias, Nax, Heinrich H. and Sornette, Didier (2015) *Super-exponential growth expectations and the global financial crisis.* Journal of Economic Dynamics and Control, 55. pp. 1-13. ISSN 0165-1889

DOI: 10.1016/j.jedc.2015.03.005

Super-Exponential Growth Expectations and the Global Financial Crisis

Matthias Leiss*, Heinrich H. Nax†, Didier Sornette‡§

January 27, 2015

We construct risk-neutral return probability distributions from S&P 500 options data over the decade 2003 to 2013, separable into pre-crisis, crisis and post-crisis regimes. The pre-crisis period is characterized by increasing realized and, especially, option-implied returns. This translates into transient unsustainable price growth that may be identified as a bubble. Granger tests detect causality running from option-implied returns to Treasury Bill yields in the pre-crisis regime with a lag of a few days, and the other way round during the post-crisis regime with much longer lags (50 to 200 days). This suggests a transition from an abnormal regime preceding the crisis to a "new normal" post-crisis. The difference between realized and option-implied returns remains roughly constant prior to the crisis but diverges in the post-crisis phase, which may be interpreted as an increase of the representative investor's risk aversion.

Keywords: financial crisis, returns, expectations, options, risk-neutral densities

JEL: D84, G01, G13, G14

*Corresponding author: ETH Zürich, Department of Social Sciences, Clausiusstrasse 50, 8092 Zürich, Switzerland. mleiss@ethz.ch. Leiss acknowledges financial support from ETH Zurich Risk Center.

†ETH Zürich, Department of Social Sciences, Clausiusstrasse 50, 8092 Zürich, Switzerland. Nax acknowledges support from ERC (Momentum).

‡ETH Zürich, Department of Management, Technology and Economics, Scheuchzerstrasse 7, CH-8092 Zürich, Switzerland

§Swiss Finance Institute, University of Geneva, 40 Blvd. du Pont d'Arve, CH-1211 Geneva, Switzerland

1. Introduction

The Global Financial Crisis of 2008 brought a sudden end to a widespread market exuberance in investors' expectations. A number of scholars and pundits had warned ex ante of the non-sustainability of certain pre-crisis economic developments, as documented by Bezemer (2011). Those who warned of the crisis identified as the common elements in their thinking the destabilizing role of uncontrolled expansion of financial assets and debt, the flow of funds, and the impact of behaviors resulting from uncertainty and bounded rationality. However, these analyses were strongly at variance with the widespread belief in the "Great Moderation" (Stock and Watson, 2003) and in the beneficial and stabilizing properties of financial derivatives markets by their supposed virtue of dispersing risk globally (Summers et al., 1999; Greenspan, 2005). In hindsight, it became clear to everyone that it was a grave mistake to ignore issues related to systemic coupling and resulting cascade risks (Bartram et al., 2009; Hellwig, 2009). But could we do better in the future and identify unsustainable market exuberance ex ante, to diagnose stress in the system in real time before a crisis starts?

The present article offers a new perspective on identifying growing risk by focussing on growth expectations embodied in financial option markets. We analyze data from the decade around the Global Financial Crisis of 2008 over the period from 2003 to 2013.[1] We retrieve the full risk-neutral probability measure of implied returns and analyze its characteristics over the course of the last decade. Applying a change point detection method (Killick et al., 2012), we endogenously identify the beginning and end of the Global Financial Crisis as indicated by the options data. We consistently identify the beginning and end of the Crisis to be June 2007 and May 2009, which is in agreement with the timeline given by the Federal Reserve Bank of St. Louis (2009).[2]

The resulting pre-crisis, crisis and post-crisis regimes differ from each other in several important aspects. First, during the pre-crisis period, but not in the crisis and post-crisis periods, we identify a continuing increase of S&P 500 expected returns. This corresponds to super-exponential growth expectations of the price. By contrast, regular expectation regimes prevail in the crisis and post-crisis periods. Second, the difference between realized and option-implied returns remains roughly constant prior to the crisis but diverges in the post-crisis phase. This phenomenon may be interpreted as an increase of the representative investor's risk aversion. Third, Granger-causality tests show that changes of option-implied returns Granger-cause changes of Treasury Bill yields with a lag of few days in the pre-crisis period, while the reverse is true at lags of 50 to 200 days in the post-crisis period. This role reversal suggests that Fed policy was responding to, rather than leading, the financial market development during the pre-crisis period, but that the economy returned to a "new normal" regime post-crisis.

[1] Related existing work has considered data from pre-crisis (Figlewski, 2010) and crisis (Birru and Figlewski, 2012).

[2] See section 3.2 for more details on market and policy events marking the Global Financial Crisis of 2008.

The majority of related option market studies have used option data for the evaluation of risk. An early contribution to this strand of work is Aït-Sahalia and Lo (2000) who proposed a nonparametric risk management approach based on a value at risk computation with option-implied state-price densities. Another popular measure of option-implied volatility is the Volatility Index (VIX), which is constructed out of options on the S&P 500 stock index and is meant to represent the market's expectation of stock market volatility over the next 30 days (Chicago Board Options Exchange, 2009). Bollerslev and Todorov (2011) extended the VIX framework to an "investor fears index" by estimating jump tail risk for the left and right tail separately. Bali et al. (2011) define a general option-implied measure of riskiness taking into account an investor's utility and wealth leading to asset allocation implications. What sets our work apart is the focus on identifying the long and often slow build-up of risk during an irrationally exuberant market that typically precedes a crisis.

Inverting the same logic, scholars have used option price data to estimate the risk attitude of the representative investor as well as its changes. These studies, however, typically impose stationarity in one way or another. Jackwerth (2000), for example, empirically derives risk aversion functions from option prices and realized returns on the S&P 500 index around the crash of 1987 by assuming a constant return probability distribution. In a similar way, Rosenberg and Engle (2002) analyze the S&P 500 over four years in the early 1990s by fitting a stochastic volatility model with constant parameters. Bliss and Panigirtzoglou (2004), working with data for the FTSE 100 and S&P 500, propose another approach that assumes stationarity in the risk aversion functions. Whereas imposing stationarity is already questionable in "normal" times, it is certainly hard to justify for a time period covering markedly different regimes as around the Global Financial Crisis of 2008. We therefore proceed differently and merely relate return expectations implicit in option prices to market developments, in particular to the S&P 500 stock index and yields on Treasury Bills. We use the resulting data trends explicitly to identify the pre-crisis exuberance in the trends of market expectations and to make comparative statements about changing risk attitudes in the market.

The importance of market expectation trends has not escaped the attention of many researchers who focus on 'bubbles' (Galbraith, 2009; Sornette, 2003; Shiller, 2005; Soros, 2009; Kindleberger and Aliber, 2011). One of us summarizes their role as follows: "In a given financial bubble, it is the expectation of future earnings rather than present economic reality that motivates the average investor. History provides many examples of bubbles driven by unrealistic expectations of future earnings followed by crashes" (Sornette, 2014). While there is an enormous econometric literature on attempts to test whether a market is in a bubble or not, to our knowledge our approach is the first trying to do so by measuring and evaluating the market's expectations directly.[3]

[3] For the econometric literature regarding assessments as to whether a market is in a bubble or not see Stiglitz (1990) (and the corresponding special issue of the Journal of Economic Perspectives), Bhattacharya and Yu (2008) (and the corresponding special issue of the Review of Financial Studies), as well as Camerer (1989), Scheinkman and Xing (2003), Jarrow et al. (2011), Evanoff et al. (2012), Lleo and

This paper is structured as follows. Section 2 details the estimation of the risk-neutral return probability distributions, the identification of regime change points, and the causality tests regarding market returns and expectations. Section 3 summarizes our findings, in particular the evidence concerning pre-crisis growth of expected returns resulting in super-exponential price growth. Section 4 concludes with a discussion of our findings.

2. Materials and Methods

2.1. Estimating risk-neutral densities

Inferring information from option exchanges is guided by the fundamental theorem of asset pricing stating that, in a complete market, an asset price is the discounted expected value of future payoffs under the unique risk-neutral measure (see e.g. Delbaen and Schachermayer, 1994). Denoting that measure by $\mathbb{Q}$ and the risk-neutral density by f, respectively, the current price C_0 of a standard European call option on a stock with price at maturity S_T and strike K can therefore be expressed as

$$C_0(K) = \mathrm{e}^{-r_f T}\ \mathbb{E}_0^{\mathbb{Q}}\left[\max(S_T - K, 0)\right] = \mathrm{e}^{-r_f T}\ \int_K^{\infty} (S_T - K) f(S_T) \mathrm{d}S_T, \tag{1}$$

where r_f is the risk-free rate and T the time to maturity. From this equation, we would like to extract the density $f(S_T)$, as it reflects the representative investor's expectation of the future price under risk-neutrality. Since all quantities but the density are observable, inverting equation (1) for $f(S_T)$ becomes a numerical task.

Several methods for inverting have been proposed, of which Jackwerth (2004) provides an excellent review. In this study, we employ a method by Figlewski (2010) that is essentially model-free and combines standard smoothing techniques in implied-volatility space and a new method of completing the density with appropriate tails. Tails are added using the theory of Generalized Extreme Value distributions, which are capable of characterizing very different behaviors of extreme events.[4] This method cleverly combines mid-prices of call and put options by only taking into account data from at-the-money and out-of-the-money regions, thus recovering non-standard features of risk-neutral densities such as bimodality, fat tails, and general asymmetry.

Our analysis covers fundamentally different market regimes around the Global Financial Crisis. A largely nonparametric approach, rather than a parametric one, seems therefore appropriate, because an important question that we shall ask is whether and how distributions actually changed from one regime to the next. We follow Figlewski's method in

Ziemba (2012), Anderson et al. (2013), Phillips et al. (2013), Hüsler et al. (2013).

[4] As Birru and Figlewski (2012) note, the theoretically correct extreme value distribution class is the Generalized Pareto Distribution (GPD) because estimating beyond the range of observable strikes corresponds to the peak-over-threshold method. For our purposes, both approaches are known to lead to equivalent results.

most steps, and additionally weight points by open interest when interpolating in implied-volatility space – a proxy of the information content of individual sampling points permitted by our data. We give a more detailed review of the method in appendix A.

2.2. Data

We use end-of-day data for standard European call and put options on the S&P 500 stock index provided by Stricknet[5] for a period from January 1st, 2003 to October 23rd, 2013. The raw data includes bid and ask quotes as well as open interest across various maturities. For this study, we focus on option contracts with quarterly expiration dates, which usually fall on the Saturday following the third Friday in March, June, September and December, respectively. Closing prices of the index, dividend yields and interest rates of the 3-month Treasury Bill as a proxy of the risk-free rate are extracted from Thomson Reuters Datastream.

We apply the following filter criteria as in Figlewski (2010). We ignore quotes with bids below \$0.50 and those that are larger than \$20.00 in the money, as such bids exhibit very large spreads. Data points for which the midprice violates no-arbitrage conditions are also excluded. Options with time to maturity of less than 14 calendar days are discarded, as the relevant strike ranges shrink to smaller and smaller lengths resulting in a strong peaking of the density.[6] We are thus left with data for 2,311 observations over the whole time period and estimate risk-neutral densities and implied quantities for each of these days.

2.3. Subperiod classification

As the Global Financial Crisis had a profound and lasting impact on option-implied quantities, it is informative for the sake of comparison to perform analyses to subperiods associated with regimes classifiable as pre-crisis, crisis and post-crisis. Rather than defining the relevant subperiods with historical dates, we follow an endogenous segmentation approach for identifying changes in the statistical properties of the risk-neutral densities. Let us assume we have an ordered sequence of data $x_{1:n} = (x_1, x_2, ..., x_n)$ of length n, e.g. daily values of a moment or tail shape parameter of the risk-neutral densities over n days. A change point occurs if there exists a time $1 \leq k < n$ such that the mean of set $\{x_1, ..., x_k\}$ is statistically different from the mean of set $\{x_{k+1}, ..., x_n\}$ (Killick et al., 2012). As a sequence of data may also have multiple change points, various frameworks to search for them have been developed. The binary segmentation algorithm by Scott and Knott (1974) is arguably the most established detection method of this kind. It starts by identifying a single change point in a data sequence, proceeds iteratively on the two segments before and after the detected change and stops if no further change point is found.

[5] The data is accessible via stricknet.com, where it can be purchased retrospectively.

[6] Figlewski (2010) points out that rollovers of hedge positions into later maturities around contract expirations may lead to badly behaved risk-neutral density estimates.

As in the case of estimating risk-neutral densities, we refrain from making assumptions regarding the underlying process that generates the densities and choose a nonparametric approach. We employ the numerical implementation of the binary segmentation algorithm by Killick et al. (2012) with the cumulative sum test statistic (CUSUM) proposed by Page (1954) to search for at most two change points. The idea is that the cumulative sum, $S(t) := \sum_{i=1}^{t} x_i$, $1 \leq t < n$, will have different slopes before and after the change point. As opposed to moving averages, using cumulative sums allows rapid detection of both small and large changes. We state the mathematical formulation of the test statistic in appendix B.[7]

2.4. Determining lag-lead structures

Option-implied quantities may be seen as expectations of the (representative) investor under $\mathbb{Q}$. A popular question in the context of self-referential financial markets is whether expectations drive prices or vice versa. To get a feeling of the causality, we analyze the lag-lead structure between the time series based on the classical method due to Granger (1969). Informally, 'Granger causality' means that the knowledge of one quantity is useful in forecasting another. Formally, given two time series X_t and Y_t, we test whether Y_t Granger-causes X_t at lag m as follows. We first estimate the univariate autoregression

$$X_t = \sum_{j=1}^{m} a_j X_{t-j} + \varepsilon_t, \tag{2}$$

where ε_t is an uncorrelated white-noise series. We then estimate the augmented model with lagged variables

$$X_t = \sum_{j=1}^{m} b_j X_{t-j} + \sum_{j=1}^{m} c_j Y_{t-j} + \nu_t, \tag{3}$$

where ν_t is another uncorrelated white-noise series. An F-test shows if the lagged variables collectively add explanatory power. The null hypothesis "Y_t does not Granger cause X_t" is that the unrestricted model (3) does not provide a significantly better fit than the restricted model (2). It is rejected if the coefficients $\{c_j, j = 1...m\}$ are statistically different from zero as a group. Since the model is only defined for stationary time series, we will test for Granger causality with standardized incremental time series in identified subperiods as described in section 2.3.

[7] Interested readers may consult Brodsky and Darkhovsky (1993) as well as Csörgö and Horváth (1997) for a deeper discussion of theory, applications, and potential pitfalls of these methods.

3. Results

3.1. First-to-fourth return moment analyses

We start by analyzing the moments and tail shape parameters of the option-implied risk-neutral densities over the whole period (see Figure 1). For comparability, we rescale the price densities by the S&P 500 index level S_t, i.e. assess $f(S_T/S_t)$ instead of $f(S_T)$.[8] In general, we recover similar values to the ones found by Figlewski (2010) over the period 1996 to 2008. The annualized option-implied log-returns of the S&P 500 stock index excluding dividends are defined as

$$r_t = \frac{1}{T-t}\int_0^\infty \log\left(\frac{S_T}{S_t}\right) f(S_T)\mathrm{d}S_T. \tag{4}$$

They are on average negative with a mean value of -3%, and exhibit strong fluctuations with a standard deviation of 4%. This surprising finding may be explained by the impact of the Global Financial Crisis and by risk aversion of investors as explained below. The annualized second moment, also called risk-neutral volatility, is on average 20% (standard deviation of 8%). During the crisis from June 22nd, 2007 to May 4th, 2009, we observe an increase in risk-neutral volatility to $29 \pm 12\%$.

A skewness of -1.5 ± 0.9 and excess kurtosis of 10 ± 12 indicate strong deviations from log-normality, albeit subject to large fluctuations.[9] During the crisis, we measure a third (-0.9 ± 0.3) and fourth moment (4.4 ± 1.6) of the risk-neutral densities closer to those of a log-normal distribution than before or after the crisis. Birru and Figlewski (2012) find a similar dynamic using intraday prices for S&P 500 Index options. For the period from September 2006 until October 2007, they report an average skewness of -1.9 and excess kurtosis of 11.9, whereas from September to November 2008 these quantities change to -0.7 and 3.5, respectively.

As the fourth moment is difficult to interpret for a strongly skewed density, one must be careful with the implication of these findings. One interpretation is that, during crisis, investors put less emphasis on rare extreme events or potential losses, that is, on fat tails or leptokurtosis, while immediate exposure through a high standard deviation (realized risk) gains importance.[10] Another interpretation of the low kurtosis and large volatility observed during the crisis regime would be in terms of the mechanical consequences of conditional estimations. The following simple example illustrates this. Suppose that the

[8] We do not go into the analysis of the first moment, which, in line with efficient markets, is equal to 1 by construction of $f(S_T/S_t)$ (up to discounting).

[9] For the sake of comparison, note that a log-normal distribution with standard deviation 20% has skewness of 0.6 and excess kurtosis of 0.7. In particular, skewness is always positive.

[10] In other words, this interpretation indicates that investors, during crisis, focus on the unfolding risk, while, during non-crisis regimes, investors worry more about possible/unlikely worst case scenarios. Related to this interpretation are hypothesis regarding human behavioral traits according to which risk-aversion versus risk-taking behaviors are modulated by levels of available attention (Gifford, Gifford).

distribution of daily returns is the sum of two Normal laws with standard deviations 3% and 20% and weights 99% and 1% respectively. This means that 99% of the returns are normally distributed with a standard deviation of 3%, and that 1% of the returns are drawn from a Gaussian distribution with a standard deviation of 20%. By construction, the unconditional excess kurtosis is non zero (27 for the above numerical example). Suppose that one observes a rare spell of large negative returns in the range of -20%. Conditional on these realizations, the estimated volatility is large, roughly 20%, while the excess kurtosis close to 0 a consequence of sampling the second Gaussian law (and Gaussian distributions have by construction zero excess kurtosis).

It is interesting to note that Jackwerth and Rubinstein (1996) reported opposite behaviors in an early derivation of the risk-neutral probability distributions of European options on the S&P 500 for the period before and after the crash of October 1987. They observed that the risk-neutral probability of a one-standard deviation loss is larger after the crash than before, while the reverse is true for higher-level standard deviation losses. The explanation is that, after the 1987 crash, option traders realized that large tail risks were incorrectly priced, and that the volatility smile was born as a result thereafter (Mackenzie, 2008).

The left tail shape parameter ξ with values of 0.03 ± 0.23 is surprisingly small: a value around zero implies that losses are distributed according to a thin tail.[11] Moreover, with -0.19 ± 0.07, the shape parameter ξ for the right tail is consistently negative indicating a distribution with compact support, that is, a finite tail for expected gains.

3.2. Regime change points

A striking feature of the time series of the moments and shape parameters is a change of regime related to the Global Financial Crisis, which is the basis of our subperiod classification. A change point analysis of the left tail shape parameter identifies the crisis period as starting from June 22nd, 2007 and ending in May 4th, 2009. As we obtain similar dates up to a few months for the change points in risk-neutral volatility, skewness and kurtosis, this identification is robust and reliable (see Table 1 for details). Indeed, the determination of the beginning of the crisis as June 2007 is in agreement with the timeline of the build-up of the financial crisis[12] (Federal Reserve Bank of St. Louis, 2009), opening the gates of loss and bankruptcy announcements. Interestingly, when applying the analysis to

[11] When positive, the tail shape parameter ξ is related to the exponent α of the asymptotic power law tail by $\alpha = 1/\xi$.

[12] (i) S&P's and Moody's Investor Services downgraded over 100 bonds backed by second-lien subprime mortgages on June 1, 2007, (ii) Bear Stearns suspended redemption of its credit strategy funds on June 7, 2007, (iii) S&P put 612 securities backed by subprime residential mortgages on credit watch, (iv) Countrywide Financial warned of "difficult conditions" on July 24, 2007, (v) American Home Mortgage Investment Corporation filed for Chapter 11 bankruptcy protection on July 31, 2007 and (vi) BNP Paribas, France's largest bank, halted redemptions on three investment funds on Aug. 9, 2007 and so on.

option-implied returns instead, we detect the onset of the crisis only on September 5th, 2008, more than a year later. This reflects a time lag of the market to fully endogenize the consequences and implication of the crisis. This is in line with the fact that most authorities (Federal Reserve, US Treasury, etc.) were downplaying the nature and severity of the crisis, whose full blown amplitude became apparent to all only with the Lehmann Brother bankruptcy.

The identification of the end of the crisis in May 2009 is confirmed by the timing of the surge of actions from the Federal Reserve and the US Treasury Department to salvage the banks and boost the economy via "quantitative easing", first implemented in the first quarter of 2009.[13] Another sign of a change of regime, which can be interpreted as the end of the crisis per se, is the strong rebound of the US stock market that started in March 2009, thus ending a strongly bearish regime characterized by a cumulative loss of more than 60% since its peak in October 2007.

Finally, note that the higher moments and tail shape parameters of the risk-neutral return densities in the post-crisis period from May 4th, 2009 to October 23, 2013 progressively recovered their pre-crisis levels.

3.3. Super-exponential return: bubble behavior before the crash

Apart from the market free fall, which was at its worst in September 2008, the second most remarkable feature of the time series of option-implied stock returns shown in Figures 1a and 2a is its regular *rise* in the years prior to the crisis. For the pre-crisis period from January 2003 to June 2007, a linear model estimates an average increase in the option-implied return of about 0.01% per trading day (p-value < 0.001, $R^2 = 0.82$, more details can be found in Table 2). As a matter of fact, this increase is also present in the realized returns, from January 2003 until October 2007, i.e. over a slightly longer period, as shown in Figure 2a. Note, however, that realized returns have a less regular behavior than the ones implied by options since the former are realized whereas the latter are expected under $\mathbb{Q}$. An appropriate smoothing such as the exponentially weighted moving average is required to reveal the trend, see Figure 2a for more details.

In the post-crisis period, in contrast, the option-implied returns exhibit less regularity, with smaller upward trends punctuated by abrupt drops. We find that option-implied returns rise on average 0.003% per trading day from May 2009 to October 2013 (p-value < 0.001). However, a coefficient of determination of $R^2 = 0.20$ suggests that this period is in fact not well-described by a linear model.

To the best of our knowledge, super-exponential price growth expectations have not previously been identified as implied by options data. This finding has several important implications that we shall now detail.

---

[13] On March 18, 2009 the Federal Reserve announced to purchase \$750 billion of mortgage-backed securities and up to \$300 billion of longer-term Treasury securities within the subsequent year, with other central banks such as the Bank of England taking similar measures.

The upward trends of both option-implied and realized returns pre-crisis signal a transient "super-exponential" behavior of the market price, here of the S&P500 index. To see this, if the average return $r(t) := \ln[p(t)/p(t-1)]$ grows, say, linearly according to $r(t) \approx r_0 + \gamma t$ as can be approximately observed in Figure 2a from 2003 to 2007, this implies $p(t) = p(t-1)e^{r_0+\gamma t}$, whose solution is $p(t) = p(0)e^{r_0 t+\gamma t^2}$. In absence of the rise of return ($\gamma = 0$), this recovers the standard exponential growth associated with the usual compounding of interests. However, as soon as $\gamma > 0$, the price is growing much faster, in this case as $\sim e^{t^2}$. Any price growth of the form $\sim e^{t^\beta}$ with $\beta > 1$ is faster than exponential and is thus referred to as "super-exponential." Consequently, if the rise of returns is faster than linear, the super-exponential acceleration of the price is even more pronounced. For instance, Hüsler et al. (2013) reported empirical evidence of the super-exponential behaviour $p(t) \sim e^{e^t}$ in controlled lab experiments (which corresponds formally to the limit $\beta \to \infty$). Corsi and Sornette (2014) presented a simple model of positive feedback between the growth of the financial sector and that of the real economy, which predicts even faster super-exponential behaviour, termed transient finite-time singularity (FTS). This dynamics can be captured approximately by the novel FTS-GARCH, which is found to achieve good fit for bubble regimes (Corsi and Sornette, 2014). The phenomenon of super-exponential price growth during a bubble can be accommodated within the framework of a rational expectation bubble (Blanchard, 1979; Blanchard and Watson, 1982), using for instance the approach of Johansen et al. (1999, 2000) (JLS model).[14] In a nutshell, these models represent crashes by jumps, whose expectations yield the crash hazard rate. Consequently, the condition of no-arbitrage translates into a proportionality between the crash hazard rate and the instantaneous conditional return: as the return increases, the crash hazard rate grows and a crash eventually breaks the price unsustainable ascension. See Sornette et al. (2013) for a recent review of many of these models.

Because super-exponential price growth constitutes a deviation from a long-term trend[15] that can only be transient, it provides a clear signature of a non-sustainable regime whose growing return at the same time embodies and feeds over-optimism and herding through various positive feedback loops. This feature is precisely what allows the association of these transient super-exponential regimes with what is usually called a "bubble" (Kaizoji and Sornette, 2009), an approach that has allowed bubble diagnostics ex-post and ex-ante (see e.g. Johansen et al., 1999; Sornette, 2003; Lin and Sornette, 2011; Sornette and Cauwels, 2014a,b).

[14] Alternative rational expectations frameworks include Sornette and Andersen (2002); Lin and Sornette (2011); Lin et al. (2014). Also related is the literature on mildly explosive bubbles (Phillips et al., 2011, 2012).

[15] Long-term exponential growth is the norm in economics, finance and demographics. This simply reflects the Gibrat law of proportional growth (Gibrat, 1931), which has an extremely broad domain of application (Yule, 1925; Simon, 1955; Saichev et al., 2009).

3.4. Dynamics of realized and option-implied returns

Realized S&P 500 and option-implied S&P 500 returns exhibit different behaviors over time (Figure 2a). Note that this difference persists even after filtering out short-term fluctuations in the realized returns.[16] During the pre-crisis period (from January 2003 to June 2007), the two grow at roughly the same rate, but the realized returns grow are approximately 8% larger than the option-implied returns. This difference can be ascribed to the "risk premium" that investors require to invest in the stock market, given their aggregate risk aversion.[17] This interpretation of the difference between the two return quantities as a risk premium, which one may literally term "realized-minus-implied risk premium", is based on the fact that the option-implied return is determined under the risk-neutral probability measure while the realized return is, by construction, unfolding under the real-world probability measure.[18] In other words, the risk-neutral world is characterized by the assumption that all investors agree on asset prices just on the basis of fair valuation. In contrast, real-world investors are in general risk-adverse and require an additional premium to accept the risks associated with their investments. During the crisis, realized returns plunged faster and deeper in negative territory than the option-implied returns, then recovered faster into positive and *growing* regimes post-crisis. Indeed, during the crisis, the realized-minus-implied risk premium surprisingly became negative.

While the option-implied returns exhibit a stable behavior punctuated by two sharp drops in 2010 and 2011 (associated with two episodes of the European sovereign debt crisis), one can observe that the realized returns have been increasing since 2009, with sharp drop interruptions, suggesting bubbly regimes diagnosed by transient super-exponential dynamics (Sornette and Cauwels, 2014b). Furthermore, the realized-minus-implied risk premium has steadily grown since 2009, reaching approximately 16% at the end of the analyzed period (October 2013), i.e. twice its pre-crisis value. This is qualitatively in agreement with other analyses (Graham and Harvey, 2013) and can be rationalized by the need for investors to be remunerated against growing uncertainties of novel kinds, such as created by unconventional policies and sluggish economic recovery.[19]

[16] Realized S&P 500 returns show more rapid fluctuations than option-implied ones, which is not surprising given that the former are realized whereas the latter are expected (under $\mathbb{Q}$). In this section we only focus on dynamics on a longer timescale, thus Figure 2a presents realized returns smoothed by an exponential weighted moving average (EWMA) of daily returns over 750 trading days. Different values or smoothing methods lead to similar outcomes.

[17] To understand variations in the risk premium in relation to the identification of different price regimes, we cannot rely on many of the important more sophisticated quantitative methods for derivation of the the risk premium, but refer to the literature discussed in the introduction. There are many avenues for promising future research to develop hybrid approaches between these more sophisticated approaches and ours which a priori allows the premium to vary freely over time.

[18] The standard definition, which usually takes the expected 10-year S&P 500 return relative to a 10-year U.S. Treasury bond yield (Fernandez, 2013; Duarte and Rosa, 2013) captures different information.

[19] An incomplete list of growing uncertainties at that time is: instabilities in the middle-East, concerns about sustainability of China's growth and issues of its on-going transitions, and many other uncertainties

3.5. Granger causality between option-implied returns and the 3-month Treasury Bill

We now examine possible Granger-causality relationships between option-implied returns and 3-month Treasury Bill yields. First note that option-implied returns and the 3-month Treasury Bill yields reveal a much weaker correlation than between realized returns and option-implied returns. A casual glance at Figure 2b suggests that their pre-crisis behaviors are similar, up to a vertical translation of approximately 3%. To see if the Fed rate policy might have been one of the drivers of the pre-crisis stock market dynamics, we perform a Granger causality test in both directions. Since a Granger test is only defined for stationary time series, we consider first differences in option-implied S&P 500 returns and 3-month Treasury Bill yields, respectively. Precisely, we define

$$SP_t = r_t - r_{t-1}, \quad TB_t = y_t - y_{t-1}. \tag{5}$$

where r_t is the option-implied return (4) and y_t is the Bill yield at trading day t. Before testing, we standardize both SP_t and TB_t, i.e. we subtract the mean and divide by the standard deviation, respectively.

There is no evidence that Federal Reserve policy has influenced risk-neutral option-implied returns over this period, as a Granger causality test fails to reject the relevant null at any lag (see Table 3 and Figure 3a). The other direction of Granger causality is more interesting, revealing Granger-causal influence of the option-implied returns on the 3-month Treasury Bill. A Granger causality test for SP_t on TB_t rejects the null for a lag of $m = 5$ trading days. This suggests that the Fed policy has been responding to, rather than leading, the development of the market expectations during the pre-crisis period. Previous works using a time-adaptive lead-lag technique had only documented that stock markets led Treasury Bills yields as well as longer term bonds yields during bubble periods (Zhou and Sornette, 2004; Guo et al., 2011). It is particularly interesting to find a Granger causality of the forward-looking expected returns, as extracted from option data, onto a backward-looking Treasury Bill yield in the pre-crisis period and the reverse thereafter. Thus, expectations were dominant in the pre-crisis period as is usually the case in efficient markets, while realized monetary policy was (and still is in significant parts) shaping expectations post-crisis (as shown in Table 3 and Figure 3b). The null of no influence is rejected for Treasury Bill yields Granger causing option-implied returns lagged by 50 to 200 days. This is coherent with the view that the Fed monetary policy, developed to catalyze economic recovery via monetary interventionism, has been the key variable influencing investors and thus options/stock markets.

Analyses of Granger causality with respect to realized returns yield no comparable results. Indeed, mutual influences with respect to Bivariate Granger tests involving the first difference time series of realized returns (with both option-implied returns and Treasury

involving other major economic players, such as Japan, India and Brazil, quantitative easing operations in the US, political will from European leaders and actions of the ECB to hold the eurozone together.

Bill yields) confirm the results that would have been expected. Both prior to and after the crisis, Treasury Bill yields Granger-cause realized returns over long time periods ($p < 0.1$ for lags of 150 and 200 trading days, respectively), whereas option-implied returns Granger-cause realized ones over short time periods ($p < 0.01$ for a lag of 5 trading days).

4. Conclusion

We have extracted risk-neutral return probability distributions from S&P 500 stock index options from 2003 to 2013. Change point analysis identifies the crisis as taking place from mid-2007 to mid-2009. The evolution of risk-neutral return probability distributions characterizing the pre-crisis, crisis and post-crisis regimes reveal a number of remarkable properties. Indeed paradoxically at first sight, the distributions of expected returns became very close to a normal distribution during the crisis period, while exhibiting strongly negative skewness and especially large kurtosis in the two other periods. This reflects that investors may care more about the risks being realized (volatility) during the crisis, while they focus on potential losses (fat left tails, negative skewness and large kurtosis) in quieter periods.

Our most noteworthy finding is the continuing increase of the option-implied average returns during the pre-crisis (from January 2003 to mid-2007), which more than parallels a corresponding increase in realized returns. While a constant average return implies standard exponential price growth, an increase of average returns translates into super-exponential price growth, which is unsustainable and therefore transient. This finding corroborates previous reports on increasing realized returns and accelerated super-exponential price trajectories, which previously have been found to be hallmarks of exuberance and bubbles preceding crashes.

Moreover, the comparison between realized and option-implied expected returns sheds new light on the development of the pre-crisis, crisis and post-crisis periods. A general feature is that realized returns adapt much faster to changes of regimes, indeed often overshooting. Interpreted as a risk premium, literally the "realized-minus-implied risk premium", these overshoots can be interpreted as transient changes in the risk perceptions of investors. We find that the realized-minus-implied risk premium was approximately 8% in the pre-crisis, and has doubled to 16% in the post-crisis period (from mid-2009 to October 2013). This increase is likely to be associated with growing uncertainties and concern with uncertainties, fostered possibly by unconventional financial and monetary policy and unexpectedly sluggish economic recovery.

Finally, our Granger causality tests demonstrate that, in the pre-crisis period, changes of option-implied returns lead changes of Treasury Bill yields with a short lag, while the reverse is true with longer lags post-crisis. In a way, the post-crisis period can thus be seen as a return to a "normal" regime in the sense of standard economic theory, according to which interest rate policy determines the price of money/borrowing, which then spills over to the real economy and the stock market. What makes it a "new normal" (El-Erian, 2011) is that zero-interest rate policies in combination with other unconventional policy actions actually dominate and bias investment opportunities. The pre-crisis reveals the opposite phenomenon in the sense that expected (and realized returns) lead the interest rate, thus in a sense "slaving" the Fed policy to the markets. It is therefore less surprising that such an abnormal period, previously referred to as the "Great Moderation" and hailed as the successful taming of recessions, was bound to end in disappointments as a bubble was built up (Sornette and Woodard, 2010; Sornette and Cauwels, 2014a).

These results make clear the existence of important time-varying dynamics in both equity and variance risk premia, as exemplified by the difference between the pre- and post-crisis periods in terms of the Granger causalities. The option-implied returns show that expectations have been changed by the 2008 crisis, and this confirms another massive change of expectations following the crash of October 1987, embodied in the appearance of the volatility smile (Mackenzie, 2008). We believe that extending our analysis to more crises will confirm the importance of accounting for changes of expectations and time-varying premia, and we will address these issues in future research.

References

Aït-Sahalia, Y. and A. W. Lo (2000). Nonparametric risk management and implied risk aversion. *Journal of Econometrics 94*(1-2), 9–51.

Bali, T. G., N. Cakici, and F. Chabi-Yo (2011). A generalized measure of riskiness. *Management Science 57*(8), 1406–1423.

Bartram, S., G. Brown, and J. Hund (2009). Estimating systemic risk in the international financial system. *Journal of Financial Economics 86*(3), 835–869.

Bezemer, D. J. (2011). The credit crisis and recession as a paradigm test. *Journal of Economics Issues 45*(1), 1–18.

Birru, J. and S. Figlewski (2012). Anatomy of a meltdown: The risk neutral density for the S&P 500 in the fall of 2008. *Journal of Financial Markets 15*(2), 151–180.

Blanchard, O. (1979). Speculative bubbles, crashes and rational expectations. *Economics Letters 3*(387-389).

Blanchard, O. and M. Watson (1982). Bubbles, rational expectations and speculative markets. In P. Wachtel (Ed.), *Crisis in Economic and Financial Structure: Bubbles, Bursts and Shocks.*

Bliss, R. R. and N. Panigirtzoglou (2004). Option-implied risk aversion estimates. *The Journal of Finance 59*(1), 407–446.

Bollerslev, T. and V. Todorov (2011). Tails, fears, and risk premia. *The Journal of Finance 66*(6), 2165–2211.

Brodsky, E. and B. Darkhovsky (1993). *Nonparametric Methods in Change Point Problems.* Mathematics and Its Applications. Springer.

Chicago Board Options Exchange (2009). The CBOE volatility index – VIX. *White Paper*.

Corsi, F. and D. Sornette (2014). Follow the money: The monetary roots of bubbles and crashes. *International Review of Financial Analysis 32*, 47–59.

Csörgö, M. and L. Horváth (1997). *Limit Theorems in Change-point Analysis.* Wiley series in probability and statistics. Wiley.

Delbaen, F. and W. Schachermayer (1994). A general version of the fundamental theorem of asset pricing. *Mathematische Annalen 300*(1), 463–520.

Duarte, F. and C. Rosa (2013). The equity risk premium: A consensus of models. *http://ssrn.com/abstract=2377504*.

El-Erian, M. (2011). Spain is not Greece and need not be Ireland. *Financial Times Retrieved 2011-08-18*, February 3.

Embrechts, P., C. Klüppelberg, and T. Mikosch (1997). *Modelling extremal events: for insurance and finance*, Volume 33. Springer.

Federal Reserve Bank of St. Louis (2009). The financial crisis: A timeline of events and policy actions. *http://timeline.stlouisfed.org/pdf/CrisisTimeline.pdf*.

Fernandez, P. (2013). The equity premium in 150 textbooks. *http://ssrn.com/abstract=1473225*.

Figlewski, S. (2010). Estimating the implied risk neutral density. In T. Bollerslev, J. Russell, and M. Watson (Eds.), *Volatility and Time Series Econometrics.* Oxford: Oxford University Press.

Galbraith, J. (2009). *The Great Crash 1929.* Mariner Books (Reprint Edition).

Gibrat, R. (1931). Les inégalités économiques: Aux inégalités des richesses, à la concentration des entreprises, aux populations des villes, aux statistiques des familles, etc., d'une loi nouvelle: La loi de l'effect proportionnel.

Gifford, S. Risk and uncertainty. *Z.J. Acs, D.B. Audretsch (eds.), Handbook of Entrepreneurship Research International Handbook Series on Entrepreneurship 5.*

Graham, J. R. and C. R. Harvey (2013). The equity risk premium in 2013. *http://ssrn.com/abstract=2206538*.

Granger, C. W. J. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica 37*(3), pp. 424–438.

Greenspan, A. (2005). *Consumer Finance.* Federal Reserve System's Fourth Annual Community Affairs Research Conference. Federal Reserve Board.

Guo, K., W.-X. Zhou, S.-W. Cheng, and D. Sornette (2011). The US stock market leads the federal funds rate and treasury bond yields. *PloS One 6*(8), e22794.

Hellwig, M. (2009). Systemic risk in the financial sector: An analysis of the subprime-mortgage financial crisis. *De Economist 157*(2), 129–207.

Hüsler, A., D. Sornette, and C. H. Hommes (2013). Super-exponential bubbles in lab experiments: evidence for anchoring over-optimistic expectations on price. *Journal of Economic Behavior & Organization 92*, 304–316.

Jackwerth, J. C. (2000). Recovering risk aversion from option prices and realized returns. *Review of Financial Studies 13*(2), 433–451.

Jackwerth, J. C. (2004). *Option-implied risk-neutral distributions and risk aversion.* Research Foundation of AIMR Charlotteville.

Jackwerth, J. C. and M. Rubinstein (1996). Recovering probability distributions from option prices. *The Journal of Finance 51*(5), 1611–1631.

Johansen, A., O. Ledoit, and D. Sornette (2000). Crashes as critical points. *International Journal of Theoretical and Applied Finance 3*(2), 219–255.

Johansen, A., D. Sornette, and O. Ledoit (1999). Predicting financial crashes using discrete scale invariance. *Journal of Risk 1*(4), 5–32.

Kaizoji, T. and D. Sornette (2009). Market bubbles and crashes. *The Encyclopedia of Quantitative Finance*, long version at http://arxiv.org/abs/0812.2449).

Killick, R., P. Fearnhead, and I. Eckley (2012). Optimal detection of changepoints with a linear computational cost. *Journal of the American Statistical Association 107*(500), 1590–1598.

Kindleberger, C. P. and R. Z. Aliber (2011). *Manias, Panics and Crashes: A History of Financial Crises.* Palgrave Macmillan; Sixth Edition.

Lin, L., R. Ren, and D. Sornette (2014). The volatility-confined lppl model: A consistent model of 'explosive' financial bubbles with mean-reversing residuals. *International Review of Financial Analysis 33*, 210–225.

Lin, L. and D. Sornette (2011). Diagnostics of rational expectation financial bubbles with stochastic mean-reverting termination times. *The European Journal of Finance DOI: 10.1080/1351847X.2011.607004*.

Lleo, B. and W. T. Ziemba (2012). Stock market crashes in 2007 – 2009: were we able to predict them? *Quantitative Finance 12*(8), 1161–1187.

Mackenzie, D. (2008). *An Engine, Not a Camera: How Financial Models Shape Markets*. The MIT Press.

Page, E. (1954). Continuous inspection schemes. *Biometrika 41*(1/2), 100–115.

Phillips, P. C. B., S.-P. Shi, and J. Yu (2012). Testing for multiple bubbles 1: Historical episodes of exuberance and collapse in the S&P 500. *Cowles Foundation Discussion Paper No. 1843*.

Phillips, P. C. B., 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.

Rosenberg, J. V. and R. F. Engle (2002). Empirical pricing kernels. *Journal of Financial Economics 64*(3), 341–372.

Saichev, A., Y. Malevergne, and D. Sornette (2009). A theory of Zipf's law and beyond. *Lecture Notes in Economics and Mathematical Systems 632*, 1–171.

Scott, A. and M. Knott (1974). A cluster analysis method for grouping means in the analysis of variance. *Biometrics 30*(3), 507–512.

Shiller, R. J. (2005). *Irrational exuberance*. Random House LLC.

Simon, H. (1955). On a class of skew distribution functions. *Biometrika 52*, 425–440.

Sornette, D. (2003). *Why stock markets crash: critical events in complex financial systems*. Princeton University Press.

Sornette, D. (2014). Physics and financial economics (1776-2013): Puzzles, Ising and agent-based models. *Rep. Prog. Phys. 77*, 062001 (28 pp.).

Sornette, D. and J. V. Andersen (2002). A nonlinear super-exponential rational model of speculative financial bubbles. *International Journal of Modern Physics C 13*(2), 171–188.

Sornette, D. and P. Cauwels (2014a). 1980-2008: The illusion of the perpetual money machine and what it bodes for the future. *Risks 2*, 103–131 (http://arxiv.org/abs/1212.2833 and `http://ssrn.com/abstract=2191509`).

Sornette, D. and P. Cauwels (2014b). Financial bubbles: Mechanism, diagnostic and state of the world (Feb. 2014). *Review of Behavioral Economics (in press)*, (http://ssrn.com/abstract=2423790).

Sornette, D. and R. Woodard (2010). Financial bubbles, real estate bubbles, derivative bubbles, and the financial and economic crisis. In *Econophysics Approaches to Large-Scale Business Data and Financial Crisis*, pp. 101–148. Springer.

Sornette, D., R. Woodard, W. Yan, and W.-X. Zhou (2013). Clarifications to questions and criticisms on the Johansen-Ledoit-Sornette bubble model. *Physica A: Statistical Mechanics and its Applications 392*(19), 4417–4428.

Soros, G. (2009). *The Crash of 2008 and What it Means: The New Paradigm for Financial Markets.* Public Affairs; Revised edition.

Stock, J. H. and M. W. Watson (2003). Has the business cycle changed and why? *NBER Macroeconomics Annual 2002, MIT Press 17*, 159–230.

Summers, L., A. Greenspan, A. Levitt, and W. Ranier (1999). *Over-the-Counter Derivatives Markets and the Commodity Exchange Act.* Report of The President's Working Group on Financial Markets.

Yule, G. U. (1925). A mathematical theory of evolution, based on the conclusions of Dr. J. C. Willis, F.R.S. *Philosophical Transactions of the Royal Society B 213*(402-410), 21–87.

Zhou, W.-X. and D. Sornette (2004). Causal slaving of the US treasury bond yield antibubble by the stock market antibubble of August 2000. *Physica A: Statistical Mechanics and its Applications 337*(3), 586–608.

Returns and distributional moments implied by S&P 500 options

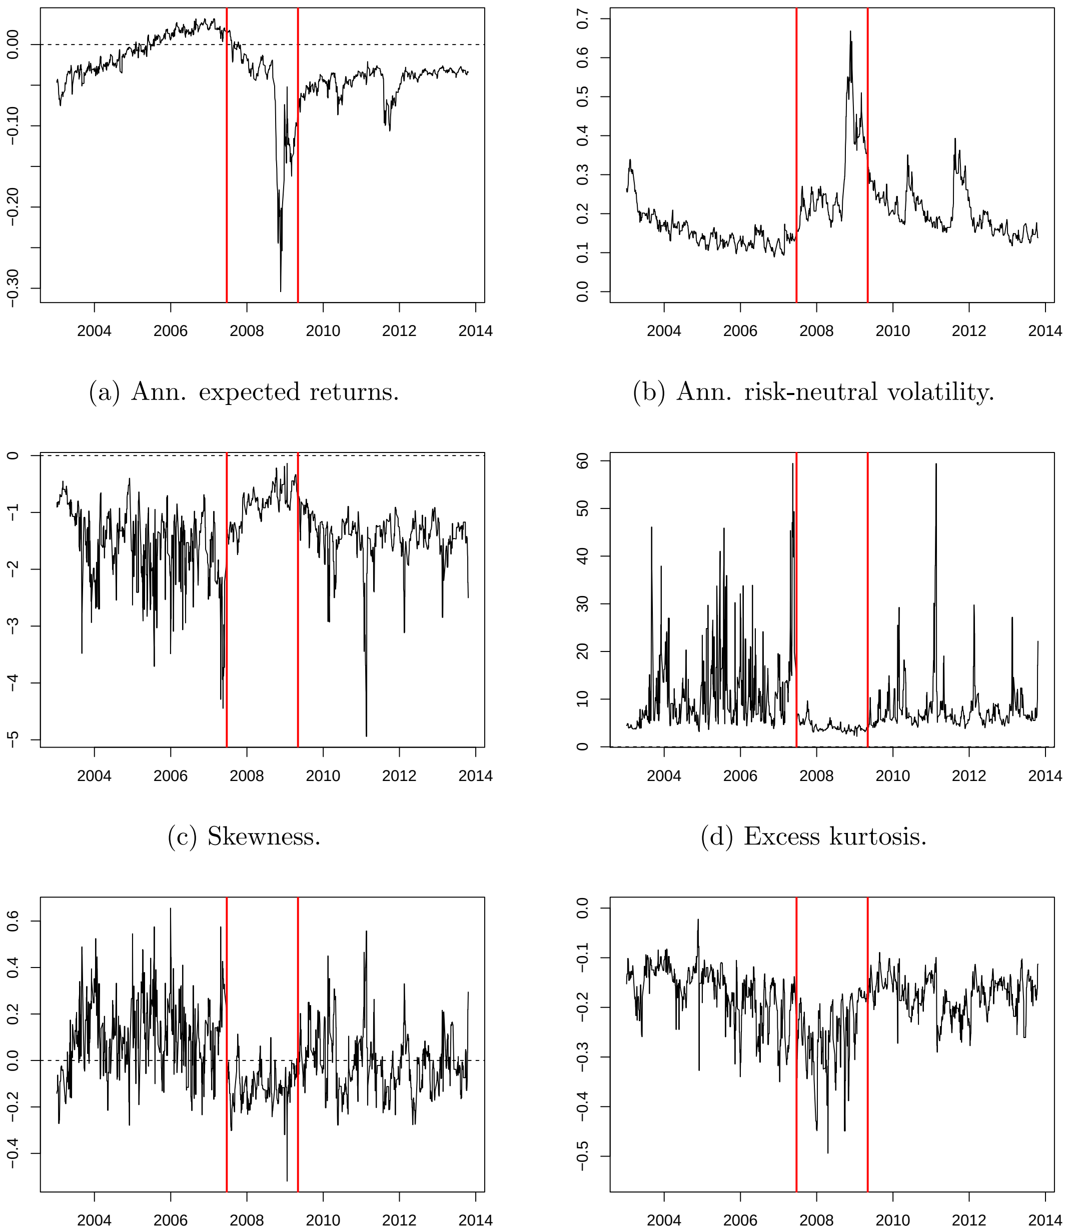

Figure 1: This figure presents returns and distributional moments implied by S&P 500 options. Structural changes around the financial crisis are identified consistently with a change point analysis of the means of the higher moments and tail shape parameters (vertical lines).

Option-implied returns vs realized returns and Treasury Bill yields

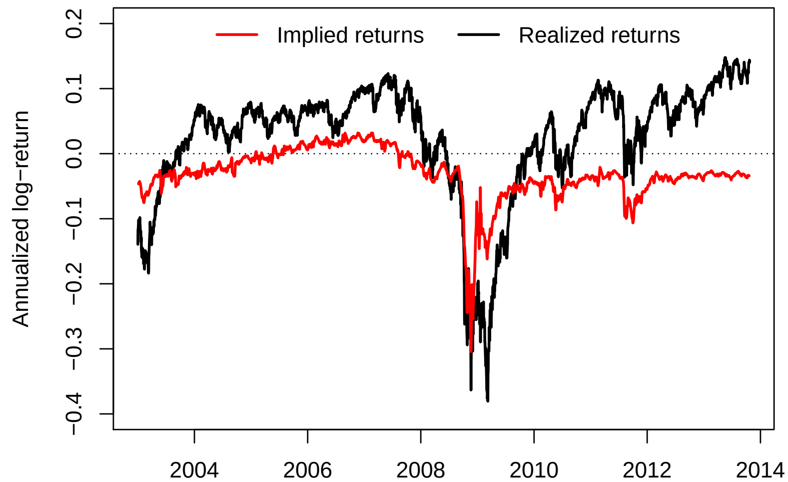

(a) Annualized realized returns and option-implied S&P 500 returns. Realized returns are calculated by exponential weighted moving average (EWMA) smoothing of daily returns over 750 trading days.

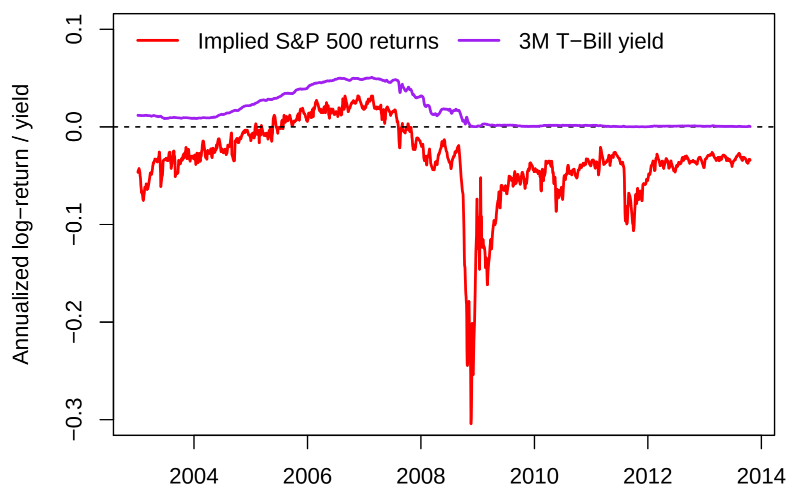

(b) 3-month Treasury Bill yields and annualized option-implied S&P 500 returns (5-day moving averages).

Figure 2: This figure presents time series of option-implied S&P 500 returns, realized returns and Treasury Bill yields over the time period 2003–2013.

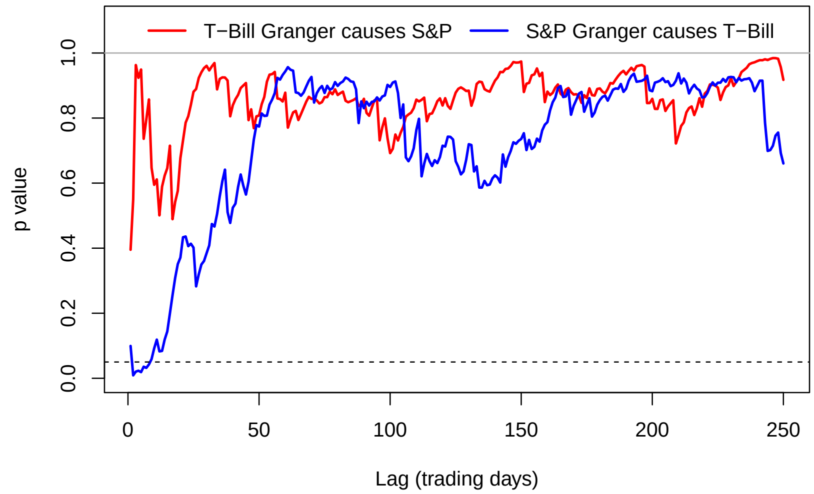

(a) Pre-crisis: January 1st, 2003 to June 22nd, 2007. No evidence for Treasury Bill yields Granger causing option-implied S&P 500 returns at any lag, but rather that option-implied S&P 500 returns Granger cause Treasury Bill yields at lags of a few trading days.

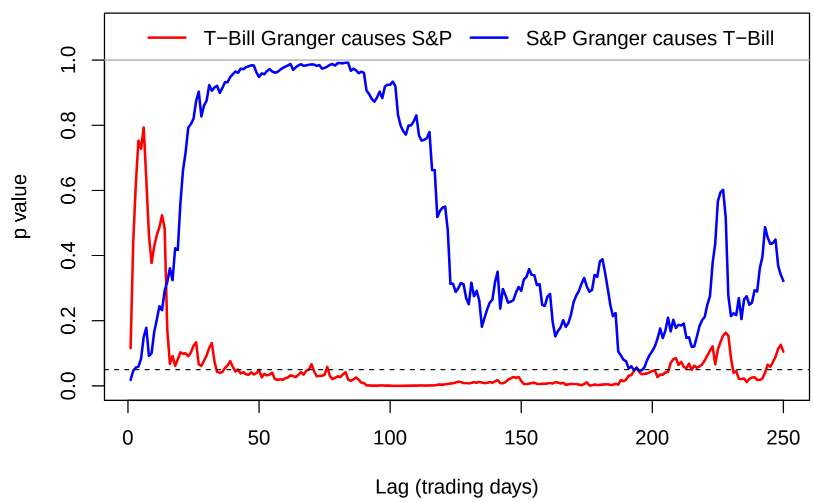

(b) Post-crisis: May 4th, 2009 to October 23rd, 2013. Treasury Bill yields Granger cause option-implied S&P 500 returns over a large range of lags.

Figure 3: Subperiod Granger causality tests on incremental changes in annualized option-implied S&P 500 returns and 3-month Treasury Bill yields. The $p = 0.05$ line is plotted as dashed black.

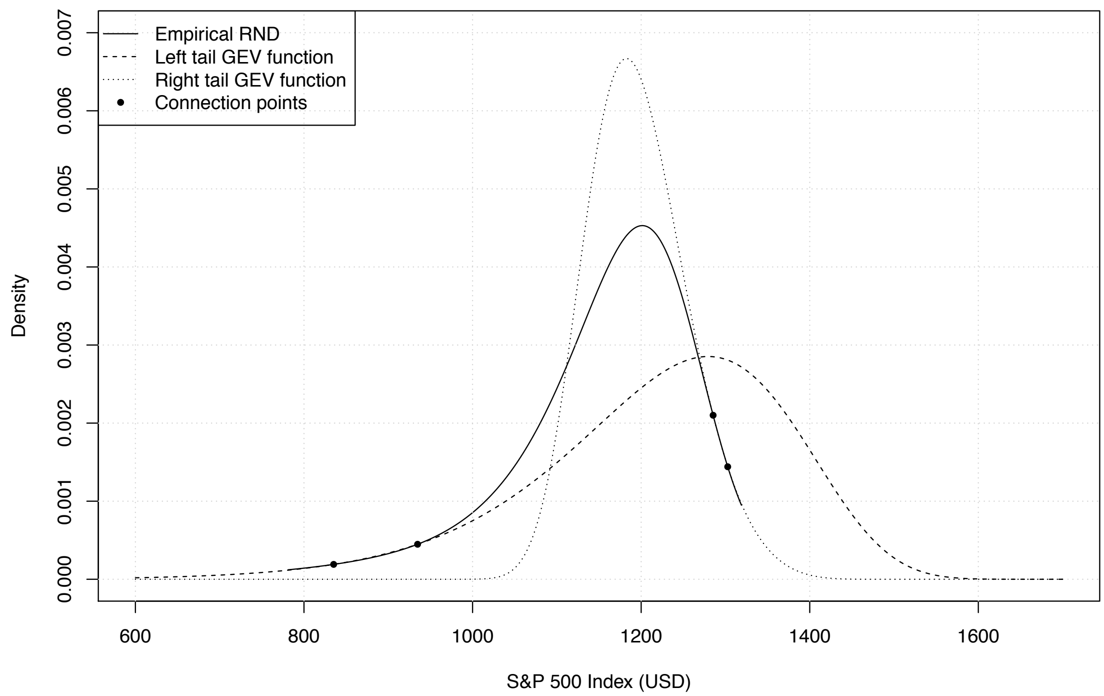

Figure 4: Risk-neutral density implied by S&P 500 options from 2010-10-06 for index levels on 2010-12-18. The empirical part is directly inferred from option quotes, whereas tails must be estimated to account for the range beyond observable strike prices. Together, they give the full risk-neutral density. The method is reviewed in section 2.1 and appendix A.

Table 1: Start and end dates of the Global Financial Crisis as identified by a change point analysis of statistical properties of option-implied risk-neutral densities. The dates found in the left tail shape parameter and higher moments identify consistently the crisis period as ca. June 2007 to ca. October 2009. Interestingly, the return time series signals the beginning only more than a year later, as September 2008. See section 2.3 for a review of the method, and 3.2 for a more detailed discussion of the results.

Variable	*Crisis start date*	*Crisis end date*
Left tail shape parameter	2007-06-22***	2009-05-04***
Right tail shape parameter	2005-08-08***	2009-01-22***
Risk-neutral volatility	2007-07-30***	2009-11-12***
Skewness	2007-06-22***	2009-10-19***
Kurtosis	2007-06-19***	*NA*[a]
Option-implied returns	2008-09-05***	2009-07-17***

Note: *p<0.1; **p<0.01; ***p<0.001

[a] No change point indicating a crisis end date found.

Table 2: Results of a linear regression of option-implied returns of the S&P 500 index on time (trading days) by sub-period. In particular, a linear model fits well the pre-crisis, indicating the regular rise of expected returns, but not the post-crisis. This translates into super-exponential price growth expectations in the pre-crisis period. Standard deviations are in parentheses.

	Option-implied returns (in percent):		
	Pre-crisis	Crisis	Post-crisis
linear coefficient	0.009^{***}	-0.043^{***}	0.003^{***}
per trading day	(0.0001)	(0.002)	(0.0002)
Constant	-4.747^{***}	2.836^{***}	-5.775^{***}
	(0.072)	(0.485)	(0.097)
Observations	942	411	958
R^2	0.820	0.520	0.196
Note:			*p<0.1; **p<0.01; ***p<0.001

Table 3: This table reports the results of a Granger-causality test of option-implied S&P 500 returns and Treasure Bill yields by sub-period. While we do not find evidence that Treasury Bill yields may have Granger-caused implied returns pre-crisis, there is Granger-influence in the other direction at a lag of 5 trading days both pre- and especially post-crisis. Notably, our test strongly suggests that post-crisis Treasury Bill yields have Granger-causal influence on option-implied returns at lags of 50 to 200 trading days.

Pre-crisis				
	S&P Granger-causes T-Bill		*T-Bill Granger-causes S&P*	
Lag	F-ratio[a]	Degrees of freedom	F-ratio[a]	Degrees of freedom
5	2.72*	5, 926	0.23	5, 926
50	0.84	50, 791	0.82	50, 791
100	0.82	100, 641	0.92	100, 641
150	0.92	150, 491	0.77	150, 491
200	0.86	200, 341	0.87	200, 341
250	0.95	250, 191	0.83	250, 191
Post-crisis				
	S&P Granger-causes T-Bill		*T-Bill Granger-causes S&P*	
Lag	F-ratio[a]	Degrees of freedom	F-ratio[a]	Degrees of freedom
5	1.95*	5, 942	0.56	5, 942
50	0.69	50, 807	1.37*	50, 807
100	0.79	100, 657	1.55**	100, 657
150	1.07	150, 507	1.32*	150, 507
200	1.16	200, 357	1.23*	200, 357
250	1.06	250, 207	1.18	250, 207

Note: *p<0.1; **p<0.01; ***p<0.001

[a] Refers to the F-test for joint significance of the lagged variables.

A. Estimating the risk-neutral density from option quotes

In this study, we estimate the option-implied risk-neutral density with a method developed by Figlewski (2010), which is based on equation (1). For completeness, we shall briefly review the method as employed in this paper, but refer the interested reader to the original document for more detail. The raw data are end-of-day bid and ask quotes of European call and put options on the S&P 500 stock market index with a chosen maturity. Very deep out of the money options exhibit spreads that are large relative to the bid, i.e. carry large noise. Due to the redundancy of calls and puts, we may discard quotes with bid prices smaller than \$0.50. In this paper, we perform the calculation with mid-prices, which by inverting the Black-Scholes model translate into implied volatilities.

In a window of ±\$20.00 around the at-the-money level, the implied volatilities of put and call options are combined as weighted averages. The weights are chosen in order to ensure a smooth transition from puts to calls by gradually blending calls into puts when going to higher strikes. Below and above that window, we only use call and put data, respectively. We then fit a fourth order polynomial in implied volatility space. Here, we deviate slightly from Figlewski (2010) because we use open interest as fitting weights. By doing so, we give more weight to data points carrying more market information. The Black-Scholes model transforms the fit in implied volatility space back to price space. The resulting density bulk is called "empirical density".

To obtain a density estimate beyond the range of observable strike prices, we must append tails to the empirical part. Figlewski (2010) proposes to add tails of the family of generalized extreme value (GEV[20]) distributions with connection conditions: a) matching value at the 2%, 5%, 92% and 95% quantile points, and b) matching probability mass in the estimated tail and empirical density. An example can be seen in Figure 4. The empirical density together with the tails give the complete risk-neutral density.

B. Change point detection

The following framework is used for significance testing in section 3.2 and Table 1. For more details, see Csörgö and Horváth (1997). Let $x_1, x_2, ..., x_n$ be independent, real-valued observations. We test the "no change point" null hypothesis,

$$H_0 : \mathbb{E}(x_1) = \mathbb{E}(x_2) = ... = \mathbb{E}(x_n), \tag{6}$$

against the "one change in mean" hypothesis,

$$H_1 : \text{there is a } k, 1 \leq k < n, \text{ such that } \mathbb{E}(x_1) = ... = \mathbb{E}(x_k) \neq \mathbb{E}(x_{k+1}) = ... = \mathbb{E}(x_n), \tag{7}$$

[20] See Embrechts et al. (1997) for a detailed theoretical discussion of GEV distributions and modeling extreme events.

using the auxiliary functions

$$A(x) := \sqrt{2 \log \log x}\ , \quad D(x) := 2 \log \log x + \frac{1}{2} \log \log \log x - \frac{1}{2} \log \pi\ . \tag{8}$$

Then, following corollary 2.1.2 and in light of remark 2.1.2. (Csörgö and Horváth, 1997, pp. 67-68), under mild regularity conditions, H_0 and for large sample sizes, one has

$$P\left(A(n) \max_k \frac{1}{\hat{\sigma}_n} \left(\frac{n}{k(n-k)} \right)^{1/2} \left| S(k) - \frac{k}{n} S(n) \right| - D(n) \leq t \right) = \exp\left(-2\mathrm{e}^{-t}\right), \tag{9}$$

where $\hat{\sigma}_n$ is the sample standard deviation and $S(t) := \sum_{i=1}^{t} x_i$ the cumulative sum of observations.