Studying dynamics in risk-neutral skewness using a Gauss-Hermite expansion on S&P 500 index options

Dennis Facius Mohit Rakyan

40968@student.hhs.se 40928@student.hhs.se

Master Thesis
Stockholm School of Economics
Tutor: Michael Halling
May 15, 2017

Abstract

Building on a new method of pricing options by modelling the underlying risk-neutral distribution with 'physicist' Hermite polynomials, we assess the properties of these distributions over time. We employ a set of S&P 500 index options ranging from 2007 to 2016. First, we provide a detailed overview of the estimation process for Gauss-Hermite risk-neutral densities. Second, we derive a simple method of directly expressing risk-neutral skewness and kurtosis from the estimated parameters of the distributions. Owing to high correlations between these moments, we focus on skewness in our analysis. We also highlight some problems associated with the tail-shape of Gauss-Hermite risk-neutral densities extracted from real-world option quotes. By careful filtering we can however remove those problematic densities. Third, we create a time-series index of skewness, which we find to be closely linked to the CBOE SKEW index. Fourth, we find that risk-neutral skewness is a mean-reverting, volatile and auto-correlated time series. However, the variation in skewness over time cannot be sufficiently modelled by autocorrelation or using returns of the underlying. Also, while we do find evidence that time-varying skewness is positively related to out-of-sample pricing errors, a large portion of the variation of risk-neutral densities remains unexplained. We conclude that while Gauss-Hermite risk-neutral densities are sufficiently well behaved to study for risk-neutral moments given appropriate filtering, our analysis suggests that they might not contain more information than other methods used in this context.

Keywords: Options, Gauss-Hermite expansion, risk-neutral density, skewness

Acknowledgements

We would like to thank Michael Halling for his knowledgeable assistance and support during all stages of this thesis. All errors are our own.

Contents

List of tables

List of figures

1. Introduction

The standard approach for pricing options is the Black-Scholes model. However, its assumptions are somewhat restrictive and do not hold up to the data very well. Specifically, it has been shown that stock-returns historically have not followed the Gaussian distribution. As this is widely known, traded option prices typically do not fully adhere to the standard model. This again has been shown with Black-Scholes implied volatilities following a U-shaped pattern when plotted against strike prices (Macbeth and Merville, 1979). This so-called volatility-smile is a way of adjusting the implied distribution for fatter tails. However, over time new models for pricing options emerged to account for a non-flat implied volatility surface (IVS). Those models are less restrictive in their assumptions and can handle non-Gaussian return distributions. By calibrating these models to the prices of traded options, the underlying risk-neutral density (RND) implied by the market can be backed out.

These RNDs can be thought of a representation of the markets risk aversion and the expected probability distribution of an assets future return. Thus, studying dynamics of these implied distributions might reveal information that can be used by market participants and central banks alike. While there is only mixed evidence about how predictive RNDs are of future returns, their higher moments can still be a way of measuring investors' expectations and market sentiment. An increasingly large body of research has been focussed on the question of how much information is embedded within option prices, and over time, this branch has progressively made use of new (option pricing) methods for backing out risk-neutral densities. However, most of the literature on this topic is comparatively mature and has been concerned with pre-2007 datasets. Also, while new formal pricing models have been introduced over time, some of them have not yet been studied in this context. We thus propose to extend on the literature by exploring the feasibility of a new option pricing model to study for RND dynamics. We also employ a more recent set of S&P 500 index options (European exercise style) covering the period from January 3, 2007 to April 29, 2016.

A promising candidate to do so is the pricing model proposed in Necula, Drimus and Farkas (2016). It builds on an expansion of the Gaussian distribution based on Hermite polynomials (thus the name Gauss-Hermite- or GH-expansion) that can be fit to option data. Also, it has been shown to price options comparatively well, both in- and out-of-sample. Necula, Drimus and Farkas (2016) report average pricing errors amounting to only 10% of a comparable expansion based method as developed in Corrado and Su (1996), as well as a slightly better performance when compared to spline-based parametrizations of the IVS. This points to comparatively stable RNDs and is the reason for us to choose this model.

Building on previous research and using this new technique of backing out RNDs, we hope to answer the following questions. First, we ask if GH-RNDs are suitable for extracting higher moments implied by option prices. This relates to the question of tail behaviour of GH-RNDs. It also requires showing how to calculate measures of asymmetry and fatness of tails of these RNDs, based on the estimated coefficients of the expansion. Second, we ask how these measures evolve over time and whether they can be used to replicate some findings in the literature.

Thus, we add to the literature in the following ways. First, we provide a detailed overview of the estimation process for Gauss-Hermite RNDs and show how the optimization can be expressed as a simple least-squares problem. Due to time-constraints we introduce a simplified method of estimating the first two moments of the distributions and slightly relax restrictions on the estimated densities. We however show that our estimation still produces comparable results to those obtained in Necula, Drimus and Farkas (2016). Second, we point out some problematic tail behaviour associated with the Gauss-Hermite method. Still, we find that careful filtering can ensure realistic RND estimates. Third, we derive simple formulas for skewness and kurtosis based on the estimated expansion coefficients, which we then use to build time-series indices of those moments. Interestingly, our skewness index is closely tied to the CBOE SKEW index, which however is calculated very differently. Fourth, we use these measures to study the dynamics of Gauss-Hermite risk-neutral distributions implied by options on the S&P 500. We find that changes in risk-neutral skewness and kurtosis are highly correlated, and for this reason we focus solely on the skewness index in the further analysis (as is often done in the literature). We also find that while the skewness index is a highly auto-correlated series, changes in the index appear to be mostly random and shocks cannot be consistently attributed to returns of the underlying. Still, skewness is related to VIX levels as in times of extremely high implied volatility, skewness is small. Also, our measure of skewness is especially volatile during times of high volatility in the underlying. This is important, as we show that time-varying skewness explains some part of the variability in the out-of-sample pricing performance of the GH-expansion. However, a large part of RND dynamics cannot be accounted to either changes in the level of implied volatility or risk-neutral skewness. While our findings mostly confirm earlier studies that have been conducted on the IVS, they also imply that GH-RNDs do not contain more information than RNDs estimated with other methods do. This assessment is however limited, as we restrict our analysis to rather simple tests. Also, we do not show if other results in the literature are replicable, e.g. if the GH-RNDs predict actual return distributions.

The rest of this paper is structured as follows: We first provide a comprehensive review of the existing literature to motivate our topic in *Chapter 2*. In *Chapter 3*, we explore one well-studied and one comparatively new method (Gauss-Hermite expansion) of backing out RNDs. We also show that skewness and kurtosis of the RND can be expressed as a linear combination of its estimated expansion coefficients, and we propose a possible definition for time-series indices of said moments. In the empirical section (*Chapter 4*), we introduce our dataset and proceed with the estimation of RNDs. We carefully compare our pricing performance with that of the relevant reference literature to ensure the validity of our optimization routine, given the simplified model employed in our paper. A short discussion about the tail-behaviour of GH-RNDs follows and we introduce some additional filters to ensure well-behaved densities. We proceed by studying the evolution of risk-neutral skewness over time, both graphically and employing time-series analysis models and regressions. In *Chapter 5,* we provide a critical assessment of our results, their implications and links to the literature, and identify topics for future research. Lastly, *Chapter 6* concludes.

2. Literature review

2.1 Option pricing

Options are derivative contracts offering the right to buy (call option) or sell (put option) an asset (underlying) at certain agreed upon prices (strike or exercise price) either at a specific date (maturity) or during a certain interval of time. Options are often traded as standardized contracts in regulated markets. Thus, their prices are set by the laws of supply and demand. However, as derivatives, their prices will also be tied to characteristics of the underlying. This and the uncertain payoffs of option contracts makes their valuation not straightforward. In the following, we highlight the evolution of important pricing models (for European exercise styles).

The first paper to evaluate stock options using mathematical methods is credited to Bachelier (1900), wherein he applied a Brownian motion to model a stochastic process, accounting for random movements in the underlying. However, by assuming the options value to rise in proportion to the square root of the time to maturity, this model significantly overpriced long-term option contracts. To obtain more realistic prices, Samuelson (1965) expanded on Bachelier's idea by introducing a geometric Brownian motion instead. However, owing to economic arguments (diversification of risk and no-arbitrage arguments), Samuelson and Merton (1969) refuted the above models.

These ideas were further revamped by Black and Scholes (1973) and Merton (1973), who pioneered the risk neutral valuation framework. The widely acknowledged and used Black-Scholes-Merton (BSM) formula gives the price of a vanilla option under a series of assumptions that are necessary to ensure the condition of no arbitrage (constant risk-free rate, no transaction costs, borrowing at the risk-free rate and no constraints on short-selling). While these assumptions do not deviate from reality by much, the assumption of a constant volatility and a Gaussian distribution however are more problematic. Several early empirical studies including Macbeth and Merville (1979) and Hull and White (1987) provided a critique of the BSM formula which misprices deep in-the money and deep out-of-the money options. These irregularities can be observed by plotting BSM-implied volatilities (IV), defined as the volatility input in the BSM model that makes the theoretical price equal to that of the traded option, against strike prices over a fixed maturity. The result is a 'U' shape within the curve, often known as a volatility smile. This smile also varies across different maturities and periods (e.g. Homescu, 2011). At any given day, the different IV's over a range of strike prices and times to maturities are referred to as an implied volatility surface (IVS). Empirically, this surface has been quite uneven, while the assumptions of the original BSM model would require

a flat surface. Thus, deviations from the BSM model again can be attributed to the restrictive assumptions of the model, especially that of log-normally distributed returns. Historically however, realized log-returns exhibited leptokurtic (fat-tailed) and asymmetrical distributions (see for example Fama, 1965 or Harris and Küçüközmen, 2001).

While practitioners simply took the IV-smile as granted and learned to work around it, it was clear that the higher implied volatilities in the tails must stem from either market-wide non-constant (or non-zero) risk-aversion or from expected distributions of the underlying at maturity that exhibited fatter tails than the Gaussian distribution. Thus, implied risk-neutral distributions seemingly deviate from the assumption of a normal distribution. This resulted in academics starting to explore pricing models using processes that could explain this behaviour, or to focus on models that are able to deal with non-Gaussian distributions in closed form.

On one hand, there are so called parametric approaches, which basically try to come up with pricing processes that are consistent with observed market prices. Early studies by Cox (1976) and Rubinstein (1983) utilized and modified the BSM formula by applying more realistic assumptions to the underlying stochastic processes and partially explained the biases prevalent in the BSM model. Also, other improvements such as the jump-diffusion process (Merton, 1976), stochastic volatility approaches (Heston, 1993), and lognormal mixture models (Melick and Thomas, 1997) have been proposed. These models further evolved into a combination of the previous techniques (Andersen, Fusari and Todorov, 2015).

On the other hand, Jarrow and Rudd (1982) argued that parametric approaches often yield distributions that cannot be integrated analytically. This reliance on numerical methods of integration increases complexity. Also, while moments of empirical distributions might be well known, it is not straightforward to account for those moments using parametric approaches. These arguments gave rise to non-parametric models. So-called Edgeworth- or Gram-Charlier Type-A (GC-A) expansions (both consist of the same series but differ in the truncation of terms and thus accuracy) have been used by Jarrow and Rudd (1982) and Corrado and Su (1996) to introduce asymmetry and fat tails to RNDs. These series are based on 'probabilist' Hermite polynomials. To ensure more realistic densities, Jondeau and Rockinger (2001) and Jurczenko, Maillet, and Negrea (2004) further added a non-negativity and a martingale restriction to earlier models, respectively. Generally, these expansion-based methods have the advantage that they can approximate non-normal distributions by modifying a Gaussian distribution with said Hermite polynomials. Also, and their coefficients directly reflect the distributions' skewness and kurtosis, and integration is straightforward. Thus, they

have been widely used in the literature. However, compared to methods of interpolating the IVS or newer expansion-based models, their pricing performance is sub-par.

This was affirmed in Necula, Drimus and Farkas (2016), who pointed towards the bad convergence properties of Edgeworth and GC-A series and instead proposed the use of Gauss-Hermite polynomials for generating a modified RND. They also derived an easy to use closed-form pricing model with coefficients that can be calibrated to option data and showed that their model significantly outperforms related methods in- and out-of-sample. Since this model to our best knowledge has not been used widely in the literature, and with its interesting properties, we decided to use it in our study of risk-neutral densities. It will be interesting to see how it performs in estimating densities that are not only able to price options more accurately, but can also be interpreted in a meaningful way.

In this regard, time-series studies have been mainly focussed on simple measures like volatility, skewness and kurtosis of the implied RND. However, compared to Edgeworth and GC-A series, the estimated coefficients of the GH-expansion cannot be interpreted as easily, especially as the series is truncated at a larger number of terms. We take this as a motivation to derive some expressions that simplify the interpretation of GH-RNDs significantly, such that one can directly calculate skewness and kurtosis from the parameters of these distributions. Again, this will be especially useful when studying how RNDs evolve over time. Before doing so, we however first proceed by introducing the existing literature on time-series dynamics of the IVS and RNDs.

2.2 Dynamics of the implied volatility surface

In the implied volatility space, several methods have been suggested to account for such dynamics. Early studies by Schönbucher (1999) and Ledoit and Santa-Clara (1998) modelled dynamics of implied volatilities jointly with those of the underlying. Thus, stochastic volatility in the underlying was shown to diffuse into implied volatilities across all strikes. Other models studied how a set of underlying factors impact the shape of the IVS. The IVS was shown to be not perfectly correlated with returns of the underlying, and their dynamics are decomposed in level, slope and convexity effects, which also have been shown to be highly auto-correlated and mean-reverting (Skiadopoulos, Hodges and Clewlow, 2000 and Cont and Da Fonseca, 2002). However, the latter two components have appeared to only account for a small portion of the variation in the IVS. In risk-neutral space, these would correspond to the overall level of implied volatility, skewness and excess kurtosis.

Additional studies by Hafner (2005) and Goncalves and Guidolin (2006) have been concerned with dynamics of the IVS both in the cross-section and as time series. While Hafner (2005) again decomposed these dynamics into different risk-factors that in turn are affected by returns in the underlying and hence are stochastic, Goncalves and Guidolin (2006) in fact found that IVS movements were highly predictable and could be economically exploited (but only when assuming low transaction costs). Building on these insights, Bedendo and Hodges (2009) and Chen, Zhou and Li (2016) recommended the use of so called Kalman filters (linear and nonlinear ones, respectively) to improve IVS models. These have been shown to significantly improve pricing performance and to generate more robust forecasts.

We conclude that dynamics of the IVS are to some degree predictable, but are also affected by more than one risk-factor, as correlation of the IVS with shocks to the underlying (and implied volatility) is far from perfect. These risk-factors correspond to the slope and curvature of the implied volatility surface, and thus must be related to skewness and kurtosis of the corresponding risk-neutral densities. Not surprisingly, these measures have been widely used in the literature on that topic.

2.3 Dynamics of risk-neutral densities

While some research has been conducted about new option pricing models to account for the volatility smile and many papers have studied its time series dynamics, a different branch of literature has been busy studying these dynamics in the risk-neutral space. They also investigated the informational content hidden in option implied RNDs and how that could be useful to market participants and regulators. Basically, they ask the question whether RNDs extracted from market traded options can be used in other areas than pure option pricing. This is interesting, as options contracts are exchange traded securities with their prices depending on *future* expected payoffs. Hence, their premia must reflect a combination of market-wide expectations and risk-aversion. (These individual effects are hard to disentangle and while some effort has been made on separating them, covering this would be beyond the scope of our paper.)

Some questions that have been asked repeatedly in this context are: a) how option prices behave prior to and following market-relevant events, b) whether option prices can predict future moves in the underlying or related, whether they can accurately predict the future return distribution and its tail-risks, c) how they are related to market sentiment and d) how they can be interpreted by central banks. We highlight the most relevant findings in the following.

An early study by Bates (1991) investigated, whether option prices reflected crash fears prior to the 1987 U.S. market crash. In fact, one year before the crash, out-of-the-money puts

were found to be relatively expensive compared to calls. This was reflected in negative risk-neutral skewness. Still, Bates (1991) rejected the hypothesis of strong fears immediately prior to 'Black Monday' (October 19, 1987). Instead, they interpreted skewness as reflecting a growing demand for insurance in a market that has generally seen growing stock valuations in the prior years. Furthermore, Bates (2000) not only indicated that negative skewness persisted even after the 1987 crash, but also that this high implied market risk was not justified, given subsequent price movements. This points to the interpretation that risk-expectations in option prices in fact might lag movements in the underlying while they do not predict future price swings. Furthermore, this points to large risk-premia (basically the difference between risk-neutral and real probabilities) associated with jump risks.

Studying this in more depth, Alonso, Blanco and Rubio (2005) found that that the hypothesis that the RND predicted the future distribution of the underlying returns could not be rejected (for the Spanish market index). However, they did find long sub-periods where this relationship did not hold. Moreover, they pointed out that the RNDs significantly underestimated the right tail of the distribution. In contrast, Shiratsuka (2001) found that for the Japanese market, implied distributions contained information about future price movements, however not to a higher degree than historical distributions did. Additionally, the RNDs shape was affected by prior moves in the underlying, while in turn the shape of the RND predicted future movements in the underlying only for some sub-sample periods - RNDs can thus not be interpreted over time in a straightforward fashion. Similarly, Birru and Figlewski (2012), while acknowledging RNDs as a very detailed snapshot of investors' expectations in times of crisis, found that the RND merely reacted to movements in the underlying. This was affirmed by Gemmill and Saflekos (2000), who also pointed out that the implied RND still might reflect investor sentiment, which however appeared not to have any forecasting ability - RNDs can thus be used by the central bank or by market participants who wish to base their trading activities on divergences between their own and market expectations. Combining these insights, we see that there is very mixed evidence about the usefulness of RNDs to predict future movements in the underlying and to anticipate market crashes. Even if there is a relationship, it generally does not hold over time.

Still, regarding the question if RNDs are affected by investors sentiment, Han (2008) conducted a thorough market research using sentiment proxies. According to his findings, there is some effect of sentiment on option prices, which turned out to be especially significant for bearish sentiment (more negatively skewed RND and a steeper volatility smile). Thus, market sentiment affects options prices, which however could not be explained by models assuming

perfect-markets and rational expectations. While agreeing with the insight that sentiment affects option prices, Andreou, Kagkadis and Philip (2012) further expanded on this by distinguishing 'rational' sentiment based on economic fundamentals and irrational sentiment or 'Error in Beliefs'. They found that up until mid-1997 (corresponding to the dataset used in Han, 2008) the irrational component had been driving option prices. However, for a longer dataset, rational expectations had a significant effect on option implied skewness, while irrational sentiment did not.

Concluding from the above studies, option prices and their RNDs have some correlation to markets future expectations and thus these are often used by central banks while conducting monetary policy research. The amount of correlation, however has been a topic of contention among many researchers. Lynch and Panigirtzoglou (2008) found that most implied RNDs were stationary and persistent with expectations reverting to their long run averages except in exceptional circumstances. They argued that these expectational variables provided little to the predictive power in terms of future macroeconomic variables over their sample period. However, this does not mean that RNDs are not useful for central banks. On the contrary, in times of monetary forward guidance, an assessment of market sentiment is of high importance to monetary authorities. Also, as Jondeau and Rockinger (2000) pointed out, confidence intervals of risk-neutral densities (of currency or interest-rate options) could be used to assess if monetary policy is viewed as credible by the market. Credibility is one of the key-factors for monetary policy to be effective.

Now obviously, we cannot test for all the questions mentioned above. Instead, we focus on some basic properties of RND dynamics. We graphically study how higher-order moments of Gauss-Hermite risk-neutral densities react to or anticipate market-relevant events. Further, we focus on simple time-series statistics, such as autocorrelations. Lastly, we test the hypothesis that returns of the underlying affect these moments, and study if that has implications for option pricing and risk management.

3. Methodology

3.1 Probability measures and derivative valuation

Before introducing the methodology used in this paper, we must first define some more theoretical concepts about risk-neutral valuation and the theory of probability measures.

In the following, we follow the definitions in Gan, Ma and Xie (2014). If Σ_0 is an algebra (i.e. a collection of subsets of set S), then a measure on Σ_0 is a set function on Σ_0 that is non-negative and *countably additive*. The latter condition requires a) the set function of an empty set to satisfy $P(\emptyset) = 0$ and b) for disjoint elements of Σ_0 noted as $F_1, F_2, \ldots$ to satisfy $P(\bigcup_{n=1}^{\infty} F_n) = \sum_{n=1}^{\infty} P(F_n)$, i.e. the probability assigned to the union of all elements equals the sum of probabilities assigned to each individual element. Under these conditions, P is a probability measure if $P(S) = 1$, and then (S, Σ_0, P) is called a probability measure space.

Simply speaking, a probability measure is thus defined as a function P in probability space, which assigns real probabilities to all possible outcomes. By definition, it has unitary mass (i.e. that the sum of probabilities of all outcomes equals one), its codomain is constrained to $[0, 1]$ (for discrete distributions) and it again satisfies the property of countable additivity.

Importantly, in finance a *risk-neutral* measure is defined as the probability measure under which the discounted expected value of an asset exactly equals its market price. This argument is closely related to the fundamental theorems of asset pricing and to no-arbitrage arguments. We call this the martingale restriction, which can be written as $E[e^{-r\tau} S_{t+\tau}]_{p_{t+\tau}} = e^{-q\tau} S_t$, where $p_{t+\tau}$ is the risk-neutral measure of asset S at future time $t + \tau$ (and as expected at t), q denotes the assets dividend yield and r is the risk-free rate. From this, it follows that the price of a derivative on asset S at time t is given as the expected value of its future payoffs under the risk-neutral measure, discounted by the risk-free rate. For a European call, it is thus dependent on the payoff function $max\ (S_{t+\tau} - K, 0)$ where K is the options exercise price, and the risk neutral density function $p_{t+\tau}(S_{t+\tau})$ of the underlying S at time $t + \tau$ (maturity). The price of such an option can be expressed as

$$C_t(r, \tau, K, p_{t+\tau}) = e^{-r\tau} \int_{-\infty}^{\infty} max\ (S_{t+\tau} - K, 0)\ p_{t+\tau}(S_{t+\tau})\ dS_{t+\tau}. \qquad (1)$$

It is easy to see that this is simply the discounted future cash flow of the option, given (risk-neutral) expectations regarding the distribution of $S_{t+\tau}$. In this basic form, the exact risk-neutral measure is not yet assigned to any specific function, and no assumptions are made about the underlying return generating process. For option pricing, this is obviously crucial. As outlined above, many pricing models have been proposed, and there exist multiple ways of

approximating RNDs. It is beyond the scope of this paper to introduce each method, but a comprehensive overview is provided in Jackwerth (2004). Rather, we will focus on suitable ways of extracting RNDs from option data, which can then be studied in our time series analysis. For this, we explore two different methods as outlined below. Generally, it will be important to note that we relax the assumption of a *unique* risk neutral measure and seek to extract densities for options (SPX) with the same underlying (S&P 500), at different dates t and with different maturity dates $t + \tau$.

3.2 Non-parametrical approximation

To get a general sense about the nature of risk neutral densities, we explore a method that does not rely on any analytic density functions or assumptions about the underlying return generating process and which is widely used in the literature. As Breeden and Litzenberger (1978) pointed out, there exists an approximate relationship between the risk neutral distribution and the (second) derivative of option premia with respect to their strike prices. It can be shown that this relationship is approximately given by

$$p_{t+\tau}(S_{t+\tau} = K) = \lim_{\Delta \to 0} e^{r\tau} \frac{1}{\Delta^2} [C_t(K + \Delta, \tau) + C_t(K - \Delta, \tau) - 2C_t(K, \tau)] \qquad (2)$$

where Δ denotes the discrete step size between strikes K_i, and $p_{t+\tau}(S_{t+\tau} = K)$ is the density at exercise price K. For a detailed discussion and derivations, see Malz (2014) or Figlewski (2010). Using this method, one can approximate the risk neutral density directly from market traded options.

However, as simple as it looks, this method in its basic form has some caveats. First, it is very sensitive to its inputs and Δ, which can yield unrealistic densities, potentially producing negative probabilities. From equation (2) we see that options with a continuum of strike prices are necessary for the equality to hold. Second, the RND only extends into the tails as far as to the second smallest and largest exercise price. Still, as the literature evolved, several extensions have been developed to overcome these issues. Here, we focus on the more recent works of Figlewski (2010) and Malz (2014).

First, to address the fact that options are only traded with exercise prices in very discrete steps, options with more strikes must be interpolated. This is often done by translating option premia to implied volatility space, interpolating volatilities by fitting a spline and finally backing out interpolated option premia. Figlewski (2010) here pointed to splines of fourth degree with one knot for the interpolation and adopting a 'smoothing spline', i.e. not forcing the estimates to go through every point. A simpler method would be to use cubic splines. However,

as already pointed out in the literature, we find that this method yields unsmooth and noisy IV data, which results in badly behaved RNDs. This is due to the fact, that option prices are not always well behaved and do not always follow smooth curves (especially in the tails where the bid-ask-spread are often quite high). This results in problematic values in the (approximate) second derivative, when forcing a spline through every observation. It has been shown before, and we also find that for smaller step sizes, negative probabilities are then more often observed. Malz (2014) proposed to solve this trade-off by finding the smallest step Δ, under which RNDs for the whole dataset are still non-negative. Since the optimal Δ is positively related to the general volatility level (Malz, 2014) and since our dataset includes the financial crisis, we would have to choose a rather large Δ. This yields a high level of discreteness. Another issue related to smoothing is that implied volatilities are best estimated from out-of-the-money options. When using splines, this requires smoothing around the current level of S_t, as one must strictly cut off in-the-money options such that there are no 'double' or overlapping observations around S_t. This would often result in a small 'jump', since implied volatilities in calls and puts do not always have the same level (see Figlewski, 2010).

Second, several methods are available to address the issue of missing observations in the tail of the distribution. Figlewski (2010) proposed the use of a generalized extreme value (GEV) distribution, by fitting its parameters such that a smooth transition between the tails and the central portion of the distribution is ensured and that the probability mass equals one. This has the advantage that fat tails can be modelled. However, even if returns have historically exhibited fat tail behaviour, Malz (2014) posed the question of how useful it is to estimate tail risk by looking at the centre of the distribution. Instead, Malz (2014) recommended to assume implied volatilities in the tail to be equal to that of the option with the lowest and highest strike price, respectively. This 'clamping' procedure simply yields a lognormally shaped tail, but is also much easier to implement in practice.

For all the improvements to this method, we decided against using it in our analysis for several reasons. First, to obtain realistic RNDs, finding the correct step size of the approximation is crucial, with optimal values to ensure non-negative probabilities varying from day to day. Results are heavily dependent on the type of interpolation used. Second, tails are not straightforward to model and their meaningfulness is up to debate. Third, these density estimates have been extensively studied before, and they do not translate into any closed form option pricing model. Lastly, time-series studies often used shape measures such as skewness and kurtosis, which however are not straightforward to calculate for discrete RNDs. For these reasons, we decided to explore a different and relatively new method of extracting RNDs.

3.3 Gauss-Hermite expansion

As mentioned in *Chapter 2.1,* a large body of research has been conducted on alternative option pricing models. Of those, one branch has been building on modifications of the Gaussian distribution, using an *Edgeworth* series expansion (Jarrow and Rudd, 1982) or a *Gram–Charlier type A* expansion (Corrado and Su, 1996, 1997), both of which make use of so called *'probabilist' Hermite polynomials*. Multiple papers followed this approach and modified it over the years. Necula, Drimus and Farkas (2016) provide a comprehensive overview of the topic, but also point to deficiencies of this method. Namely, they criticise the bad convergence properties of these polynomials and instead advocate the use of *'physicist' Hermite* or *Gauss-Hermite polynomials* to model the risk-neutral measure. They also derive a closed-form option pricing formula based on these polynomials and report improved pricing performance over other (related) methods. Thus, this model has practical implications, which makes it an interesting candidate for our analysis. Also, its coefficients can be calibrated to traded options by minimizing squared pricing errors (see below).

The 'physicist' Hermite polynomials are defined recursively as

$$H_{n+1}(x) = 2xH_n(x) - 2nH_{n-1}(x) \tag{3}$$

with $H_0(x) = 1$ and $H_1(x) = 2x$ (Bronstein et al., 2013). As Necula, Drimus and Farkas (2016) point out, these polynomials converge to a fat tailed-distribution, which is an important feature in the context of option pricing. This result follows from the analysis of 'physicist' Hermite polynomials in Myller-Lebedeff (1907), who stated that *'a function $g(s)$ that, together with its first two derivatives, is continuous and finite in $-\infty \leqq s \leqq \infty$, and which vanishes for $s = \pm\infty$ of order higher than three, i.e. for which $\lim_{s\to\pm\infty} s^3 g(s) = 0$, can in this area be constructed by the series $g(s) = \sum a_n e^{-\frac{s^2}{2}} H_n(s)$, which is derived in Fourier's fashion'.* One can easily see why the modification can be based on a Gaussian distribution, which is defined as $z(s) = \frac{1}{\sqrt{2\pi}} e^{-\frac{s^2}{2}}$.

Since it is reasonable that the probability of a log-return of $\pm\infty$ is zero (i.e. the function vanishes at infinity), the GH expansion can thus be used to construct a density function that approximates the real risk neutral measure. The standardized log-return risk neutral density can then be expressed as

$$p(x) = \frac{1}{\sigma} z\left(\frac{x-\mu}{\sigma}\right) \sum_{n=0}^{\infty} a_n \, H_n\left(\frac{x-\mu}{\sigma}\right) \tag{4}$$

where $z(x)$ is the standard normal distribution, a_n are the expansion coefficients of the GH-expansion and μ is the first and σ is the second moment of the distribution (Necula, Drimus and Farkas, 2016). Note, that the GH-polynomials are being used to modify the shape of the Gaussian given the associated weights a_n. To illustrate this and to give some intuition, we plot several $z(x)H_n(x)$ in *Appendix A*. Also note, for $a_1 = 1$ and if all subsequent $a_n = 0$, the RND is simply the Gaussian distribution.

Following Necula, Drimus and Farkas (2016), we truncate the expansion after 20 terms. A higher order approximation can improve convergence, however it does so only to some limited extent. Also, the introduction of more terms is limited by the number of traded options available. Lastly, the estimation of more coefficients might prove too noisy for our time series model, especially as coefficients will naturally become extremely small for higher orders. This can also be seen in *Appendix A*, since the extreme values of Hermite polynomials are increasing with the polynomials' order *n*. Also, we find that the time needed for estimating the expansion coefficients would be increased significantly when using a higher number of terms.

Using equation (4) and truncating as above, the pricing formula of a European call with maturity $t + \tau$ may now be written as

$$C_t(S_t, K, r_t, q, \pi_{t,\tau}, \sigma_{t,\tau}, \tau) = S_t e^{\left(\pi_{t,\tau} - r + \frac{\sigma_{t,\tau}^2}{2}\right)\tau} \sum_{n=0}^{20} a_{n,t,\tau} I_n - K e^{-r\tau} \sum_{n=0}^{20} a_{n,t,\tau} J_n \quad (5)$$

where $\pi_{t,\tau}$ and $\sigma_{t,\tau}$ denote the annualized mean and volatility of the risk neutral distribution over the period τ (as a share of one year or 360 trading days) and at time t, and r is the risk-free rate. Also, I_n is given by

$$I_{n+1} = 2z(-d_1)H_n(-d_2) + 2\sigma_{t,\tau}\sqrt{\tau} I_n + 2n I_{n-1}$$

with $I_0 = N(d_1), I_1 = 2z(-d_1) + 2\sigma_{t,\tau}\sqrt{\tau} I_0$. Similarly,

$$J_{n+1} = 2z(-d_2)H_n(-d_2) + 2n J_{n-1}$$

where $J_0 = N(d_2)$ and $J_1 = 2z(-d_2)$. $N(x)$ is the cumulative standard normal distribution and

$$d_1 = \frac{log(S_t/K) + (\pi_{t,\tau} + {\sigma_{t,\tau}}^2)\tau}{\sigma_{t,\tau}\sqrt{\tau}}, d_2 = d_1 - \sigma_{t,\tau}\sqrt{\tau}.$$

For a detailed proof and derivation, please refer to the Appendix of Necula, Drimus and Farkas (2016).

It is easy to see that for $a_1 = 1$ and for all subsequent $a_n = 0$, and if $\mu = r - \frac{\sigma^2}{2}$ (i.e. enforcing that the risk-free rate equals the expected log return), equation (5) collapses to the standard Black-Scholes framework. Also note, that the pricing formula is only given for call options. Thus, to include put prices in the estimation process, one must transform them into call prices using their implied volatilities and the Black-Scholes-Framework. This is reasonable, as the risk neutral space is just a transformation of the implied volatility space, i.e. we are not interested in the actual dollar price of an option, but more in the implied volatilities. These can be expressed by calls and puts alike, without losing informational value.

To implement the model, one must first obtain estimates for risk-neutral mean and volatility. Necula, Drimus and Farkas (2016) refer to a method outlined in Bakshi, Kapadia and Madan (2003) that calculates these moments without any model-assumptions. However, as this method relies on the availability of options with a continuum of strike prices, it is extremely sensitive to the method used for interpolating prices. We were not able to identify a feasible way of doing so, such as to ensure consistently good pricing performance with the GH pricing formula. We find that the mean and volatility estimates using this method are especially sensitive to the step size of the interpolation. As an alternative, Farkas, Waelchli and Necula (2015) introduced another way of obtaining cumulants of the distribution, however this method seems to be beyond the scope of this thesis, both mathematically and given time constraints. Instead, we propose a simple way of obtaining both mean and variance, that has proven to be quite robust for us. For volatility, we choose the value that minimizes squared pricing errors across all options (of a given maturity and on a given trading day) in the standard Black-Scholes framework. For mean-estimates, we use a similar method, but with the modified Black-Scholes-framework that assumes a non-zero mean process (basically, the GH-expansion with $a_{0,\tau} = 1$ and $a_{n,\tau} = 0$ for all $n \in (1, 20)$). We believe this to be a feasible method, as the GH expansion builds on the Gaussian, such that we can approximately use the Gaussian for fitting the first two moments of the distribution to the data. Also, we find that this method yields reasonable estimates that work well in the further analysis (see *Chapter 4.2*).

It is easy to see from *Equation (5)* that the call price is linear in the GH-expansion coefficients $a_{n,\tau}$. This is helpful as we can fit those coefficients such that the sum of squared pricing errors is minimized. We can in fact severely simplify the problem at hand by employing matrix notation, such that

$$\min_{A_{t,\tau}} \left\| E_{t,\tau} - C_{t,\tau} * A_{t,\tau} \right\|^2$$

where $A_{t,\tau}$ is the vector of GH expansion coefficients associated to the expansion for options estimated at time t and with maturity $t+\tau$, $E_{t,\tau}$ denotes the column vector of observed market prices (again, where puts are transformed to calls first), and $C_{t,\tau}$ is a matrix calculated as outlined in *Appendix B*.

To obtain realistic values for $A_{t,\tau}$, one has to impose several restrictions on the optimization to ensure that the estimated density satisfies the measure properties outlined in *Chapter 3.1*. Namely, the probability mass of the RND must equal one and non-negativity of the density must be ensured. We show these restrictions in matrix notation in the *Appendix C*. Interestingly, both the constraint of unitary mass and non-negativity of the RND can be expressed in a linear form. Thus, we can use a linearly-constrained least-squares optimization algorithm for extracting estimates of $A_{t,\tau}$. We find the 'Lsqlin'-algorithm in Matlab to be quite efficient for doing so, as it produces similar estimates as 'Quadprog' (Necula, Drimus and Farkas, 2016 pointed to quadratic programming to obtain expansion coefficients) in slightly less time. Interestingly, we find that using the 'Lsqlin'-algorithm for a large consecutive number of optimizations randomly converges to non-optimal solutions, yielding extremely high pricing errors (for around 0.01% of the dataset). We noticed that some of these solutions are not reproducible and thus simply run the estimation a second time for these observations, which removes some of them. We will later show how to filter out the remaining RNDs with extreme pricing errors.

Lastly, one must also implement the martingale restriction, as pointed out in *Chapter 3.1*. For the GH-expansion with 21 terms, this is given as

$$e^{\left(\pi_{t,\tau}-(r-q)+\frac{\sigma_{t,\tau}^2}{2}\right)\tau}\sum_{n=0}^{20}a_{n,t,\tau}i^n H_n\left(-i\sigma_{t,\tau}\sqrt{\tau}\right)=1\,,i=\sqrt{-1} \tag{6}$$

(see Necula, Drimus and Farkas, 2016 for a proof). We find that in our framework of estimation, this restriction conflicts with the restriction of unitary mass. Upon implementing all constraints, we cannot get option prices to converge using the 'Lsqlin'-algorithm, i.e. the expansion consistently yields high pricing errors. However, using only the mass- and non-negativity restriction, the martingale is often violated, whereas enforcing the martingale and non-negativity yields probability masses extremely close to 1 (99.5% of the RNDs estimated later have a mass between 0.95 and 1.05, with an average mass of 0.9939). In practice, we thus decide to drop the restriction of unitary mass and only enforce non-negativity and the martingale restriction.

3.4 Higher order moments of Gauss-Hermite risk-neutral densities

In the Black-Scholes model, the underlying distribution of log-returns is assumed to follow a normal distribution. However, this assumption is not reflected in market-traded option prices. A large body of literature exists on the IVS, which is a manifestation of this anomaly. Another way of studying this behaviour is to transform the IVS into risk-neutral densities and comparing their shapes with that of a Gaussian distribution. Typically, the literature focusses on measures of asymmetry and fatness of tails. These are usually represented by the third (*skewness,* γ_1*)* and fourth (*kurtosis,* γ_2) standardized moment of a distribution. We can define these as

$$\gamma_1(x) = E\left(\frac{(X-\mu_x)^3}{\sigma^3}\right) = \int_{-\infty}^{+\infty} \frac{(X-\mu_x)^3}{\sigma^3} f(x)dx \tag{7}$$

and

$$\gamma_2(x) = E\left(\frac{(X-\mu_x)^4}{\sigma^4}\right) = \int_{-\infty}^{+\infty} \frac{(X-\mu_x)^4}{\sigma^4} f(x)dx \tag{8}$$

for a density function $f(x)$. Now, in contrast to the GC-A type expansions which are defined to have coefficients directly reflecting skewness and kurtosis, this is not the case for the GH-expansion method. Interestingly however, as we derive in *Appendix D*, for the RND represented as a Gauss-Hermite expansion, these moments can still be expressed as a simple linear combination of the (estimated) expansion coefficients and some expressions based on the polynomial's order *n*:

$$\gamma_1 = \sum_{n=0}^{\infty} a_n D_n^3, \text{where } D_n^3 = \begin{cases} (2n+1)\dfrac{2n!}{\left(\frac{n-1}{2}\right)!}, n = odd \\ 0, else \end{cases} \tag{9}$$

and

$$\gamma_2 = \sum_{n=0}^{\infty} a_n D_n^4, \text{ where } D_n^4 = \begin{cases} (4n^2+4n+3)\dfrac{n!}{\left(\frac{n}{2}\right)!}, n = even \\ 0, else \end{cases}. \tag{10}$$

These expressions are somewhat intuitive: since polynomials of even order are symmetric, they do not introduce any skewness to the RND, but contribute to kurtosis. Similarly, as polynomials of odd order are so called *odd* functions, i.e. they satisfy $-z(x)H_n(x) = z(-x)H_n(-x)$, they introduce skewness to the RND, but not kurtosis.

Lastly, we can define the *excess kurtosis* (kurtosis in excess of that of the Gaussian distribution) as $\gamma_2^e = \gamma_2 - 3$. In the following, we will use these expressions to study the shape of the GH-RND over time. For this, we create time series indices of said moments.

3.5 Time series indices of higher order moments

Let's denote the estimated skewness and kurtosis from *equations (9)* and *(10)* for RNDs of options traded at time t and with maturity $t+\tau$ by $\hat{\gamma}_{1,t}^{t+\tau}$ and $\hat{\gamma}_{2,t}^{t+\tau}$. As we defined skewness and kurtosis as normalized third and fourth moments, they do not depend on mean and variance of the distribution, and have the same scale across maturities. Thus, we create daily time series indices for skewness and kurtosis by taking the average of those moments across all available (after filtering) RNDs. We therefore define the skewness and kurtosis indices as

$$\bar{\gamma}_{1,t} = 100 - 10 * \frac{1}{N_t}\sum_{\tau}\hat{\gamma}_{1,t}^{t+\tau}, \bar{\gamma}_{2,t} = 100 + \frac{1}{N_t}\sum_{\tau}\hat{\gamma}_{2,t}^{t+\tau}. \quad (11)$$

over all available maturities $t+\tau$ at trading day t (N_t denotes the number of available maturities at time t). The method of scaling used for $\bar{\gamma}_{1,t}$ is borrowed from the CBOE SKEW index as outlined in Chicago Board Options Exchange (2010). After applying some filters (which will be discussed later), skewness is always negative, and thus higher values of $\bar{\gamma}_{1,t}$ will indicate a high *negative* skewness. The definition of $\bar{\gamma}_{2,t}$ is a modified version of the definition of the skewness, noting that our estimates of excess kurtosis are positive everywhere and larger than skewness by approximately a factor of 10. Higher values of $\bar{\gamma}_{2,t}$ indicate higher kurtosis.

This methodology of averaging observations will obviously lead to some loss of detail; a full analysis of the dataset however would be beyond the scope of this paper. Also, due to its changing composition, some noise might be introduced to the index. On the other hand, by aggregating multiple maturities, we might get smoother and more robust estimates of the moments of the RNDs for options across multiple maturities. This is the reason why we use an extended maturity band for the calculation of these indices, compared to the CBOE SKEW index that is based only on two near-term option chains. Due to our filtering techniques (see *Chapter 4.4*), the effect of the changing composition of the index would in fact become even more pronounced. Using more maturities instead, the number of dates where the composition in the index changes also increases slightly. However due to the higher number of option chains (compared to the CBOE SKEW) in the index, the effect of this rebalancing will in fact be smaller, since the weight of each individual RND is $1/N_t$. Lastly, applying a narrower maturity band, on some days the index would only be based on one RND and on others we would lose information in the longer, but still somewhat liquid option chains. Thus, we accept the introduction of noise from a time-varying composition of the index in exchange for smoothing of potentially noisy estimates of skewness and kurtosis for individual option chains, and also for the simplicity in creating a time series index that this methodology allows.

4. Empirical results

4.1 Data

Our dataset consists of SPX option quotes (closing prices) and dividend yields from January 3, 2007 to April 29, 2016. It was obtained from OptionMetrics. We also use corresponding data for S&P 500 index levels (closing prices) and the risk-free rate (one month treasury bill rate) obtained from WRDS. Lastly, we use VIX and SKEW quotes directly from the CBOE website.

We remove incomplete option quotes (where there is no implied volatility listed), options with times to maturity larger than one year, in-the-money options, as well as options with bid-prices of zero. Furthermore, we also filter the data such that the no-arbitrage condition

$$Call > max(0, Se^{-q\tau} - Ke^{-r\tau}), Put > max(0, Ke^{-r\tau} - Se^{-q\tau})$$

holds in the data (see *Expression (15)* in Farkas, Necula and Waelchli, 2015).

Additional to these general filters, to be able to interpret RND estimates in a meaningful way, we must make sure that they are representative of the 'true' RND expected by the market. Thus, an early filter (which will be expanded upon later) is added with regards to the number of available strikes at each combination of trading day and maturity. We decide to filter observations where there are less than 23 strikes available. Usually one would need a much larger sample size for a regression to be meaningful, but we are not interested in making any predictions on these regressions/coefficients. Instead, the optimization procedure is simply an exercise to fit the RND to the option data as closely as possible. Thus, any number of strikes below the number of parameters of the RND (estimated mean and volatility, and the 21 expansion coefficients) would obviously be meaningless. Observations with 23 or more strikes however should be appropriate and can be assumed to be an estimate of the 'true' RND. This makes sense, as we also have quite stringent constraints on the shape of the RND (to enforce the martingale restriction and non-negativity).

After applying these early filters, we are left with a total of 2,544,679 options with a maturity structure as shown in *Table 1*.

Table 1. Maturity structure of option data.

	Time to maturity			
	< 0.25	0.25 - 0.5	> 0.5	All
Number of options	1,741,933	406,610	396,136	2,544,679

This table shows the number of options (after filtering as outlined in *Chapter 4.1*) in the dataset and per maturity band. Time to maturity is measured as days to maturity divided by 360.

When using option data, one must make some simplifying assumptions. As Bliss and Panigirtzoglou (2002) pointed out, there are several frictions that could produce potential errors

in option data. First, errors due to data-recording and reporting can often be problematic given the less than perfect liquid nature of options. However, we believe that we can bypass this problem using several filters to clean the data. Also, rather than using raw quotes, we use reliable data from OptionMetrics. Second, liquidity-differences can potentially have some impact on pricing. However, we decide not to remove non-traded options for two reasons: to have a larger dataset to work with and for the fact that these options still reflect some market participants' willingness to trade at the given bid- and ask prices Also, we use midpoint prices. Third, infrequent trading within some options poses a problem of non-synchronicity, i.e. the daily closing price reported may not reflect current 'trading' prices for all strikes. However, in a large market with multiple market makers, and using best bid and ask-prices across several exchanges, we can rule out heavy pricing errors (except within high volatility states of the market). Still, with trading at CBOE ending 15:15 Central Standard Time (CST), we must assume that the day-end data for the S&P 500 index levels (and other variables) is recorded at that time, too.

4.2 In-sample pricing performance

We use the method outlined in *Chapter 3.3* to extract the parameters of the risk neutral densities. Filtering as above, we obtain estimates for 29,893 RNDs (each one corresponding to an individual combination of one of 2,348 trading days and 388 maturity dates). For each of these estimates, we have a sufficient number (23 or more) of options available at different strike prices. To measure pricing performance, the mean absolute pricing error (*MAE*) is given as

$$MAE = \frac{1}{N}\sum_{i=1}^{N} |c_i - \hat{c}_i|$$

where N is the number of options priced by a given RND (or alternatively the number of options in a given subset of the data) and c_i and $\hat{c}_i$ are the real (calculated from IV) and estimated prices of a call option, given the previously estimated corresponding RNDs.

However, since the definition of *MAE* does not consider differing general price levels of options with different times to maturity and time varying implied volatilities, we additionally propose to measure the pricing performance based on the model's relative absolute pricing error (RPE), which we define as

$$RPE = \frac{\sum_{i=1}^{N} |c_i - \hat{c}_i|}{\sum_{i=1}^{N} c_i}.$$

Now, in-sample pricing errors measure how well RNDs approximate real option prices. We report these in *Table 2*. For the whole dataset (after filtering in *3.1*), we find a *MAE* of 0.2. Since the maturity structure of our dataset (see *Table 1*) is different than that of Necula, Drimus and Farkas (2016), we must however compare pricing performance of individual maturity bands. We find *MAE* for short term ($\tau < 0.25$) options of 0.15, which is slightly larger than that of Necula, Drimus and Farkas (2016), who report a corresponding *MAE* of 0.11. Note, that they do not calculate relative pricing errors. Lastly, we notice that *MAEs* are higher for longer term options, but *RPEs* are stable over τ, which is due to higher premia associated with large τ.

Table 2. In-sample pricing errors of the GH model.

	Time to maturity			
	< 0.25	0.25 - 0.5	0.5 - 1	All*
MAE (USD)	0.1524	0.3626	0.2313	0.1983
	(0.0976)	(0.1456)	(0.2288)	(0.1257)
RPE	0.12%	0.16%	0.09%	0.12%
	(0.07%)	(0.06%)	(0.09%)	(0.07%)
No. of RNDs	15,339	6,501	8,053	29,893
	(15,332)	(6,497)	(8,047)	(29,876)
Average No. of	113.56	62.55	49.19	85.13
strikes per RND	(113.57)	(62.55)	(49.20)	(85.14)

Average mean absolute pricing errors of the GH expansion (after filtering as outlined in *Chapter 4.1*) over the whole dataset, whereby we weight each RND by its number of available strikes. We also report (simple) average RPEs over all observations/RNDs. Lastly, we report the number of estimated RNDs and the average number of available strikes per RND in each maturity band. Numbers in parentheses refer to results after filtering the 17 observations where the GH expansion was not able to converge to option prices and yielded a relative pricing error of more than 2.5%.

While the pricing performance on average is reasonable, we also find some extremely high errors (*MAE* well above $100), and thus instances where the GH expansion is not able to approximate option prices. These correspond not only to dates with high volatility in the markets, but also appear randomly. This can be due to some errors in the option data (as discussed above), but also due to the fact that we observe non-monotonic call prices (i.e. prices of call options on a given date/maturity combination not strictly falling for increasing exercise prices) on dates with the highest pricing errors. This can yield high errors in any model that enforces a strict probability measure, as from a pricing perspective this behaviour can only result from negative parts in the RND. In practice however, it might simply result from large differences in the IVS of call and put options, which we combine around S_t (to remove in-the money options), potentially yielding a 'jump' in prices. Since we will filter out RNDs with relative pricing errors above 2.5% later (*Chapter 4.4*), *Table 2* additionally reports pricing errors of the remaining RNDs after this filtering. Note that this only removes 17 RNDs out of

a total of 29,893 RNDs. Upon filtering these, our *MAE*s are slightly lower than those reported in Necula, Drimus and Farkas (2016) (but we do not enforce a unitary probability mass).

These results show that our implementation of the GH-expansion holds up to the literature in terms of pricing performance. This is especially important, since we use a different method to obtain estimates for mean and volatility (see a short discussion below). We also do not filter for volatility-adjusted moneyness (see Necula, Drimus and Farkas, 2016) to have more available option quotes. Judging from the pricing performance, we are confident that we can consider the estimated densities as a good approximation of the real risk neutral measures and use them in our analysis.

Concerning our simplified estimates of mean and volatility, we need to make sure that they do not induce high errors to our model in case the RNDs deviate from the Gaussian distribution. Thus, we check if there is a connection between pricing performance and the 'share' of RND that is explained by the Gaussian distribution (i.e. estimates for a_0). Naturally, since we estimate the first two moments such that pricing errors are minimized under a normal distribution, the performance of the GH-expansion might be negatively affected when the RND deviates heavily from the Gaussian distribution. To check this hypothesis, we regress the obtained *RPE* (after filtering as above) on the absolute difference of a_0 and 1 (an a_0 of 1 simply produces the Gaussian), i.e.

$$RPE = \beta|a_0 - 1| + \varepsilon$$

for which we find the p-value of β to be 0 and a R^2 of 0.34. Hence, as expected, there is a penalty associated with deviations from normality. This again might be attributed to our method of fitting mean and volatility. However, it could also be that in high volatility- and high uncertainty states of the market (where deviations from the Gaussian are the largest) prices are generally not well behaved and thus errors are naturally large. We further find evidence of this when regressing $|a_0 - 1|$ on the estimated volatility (which by design is closely linked to VIX), which again yields a p-value of the coefficient as 0 and a R^2 of 0.56. Also, we find that there is a relationship between RPE and the estimated volatility. We recognize that due to our model-bound estimation of mean and volatility, we can hardly distinguish between these individual effects. To check whether our estimators have a significantly negative effect on the performance of the GH expansion, one would have to rigorously compare it with the performance when using other, model-independent estimators, for example the ones given in Bakshi, Kapadia and Madan (2003). Again, due to time constraints, we refrain from doing so and point to the fact

that our pricing performance generally holds up to that reported in Necula, Drimus and Farkas (2016).

Figure 1. Evolution of daily average RPE from January 3, 2007 to April 29, 2016.

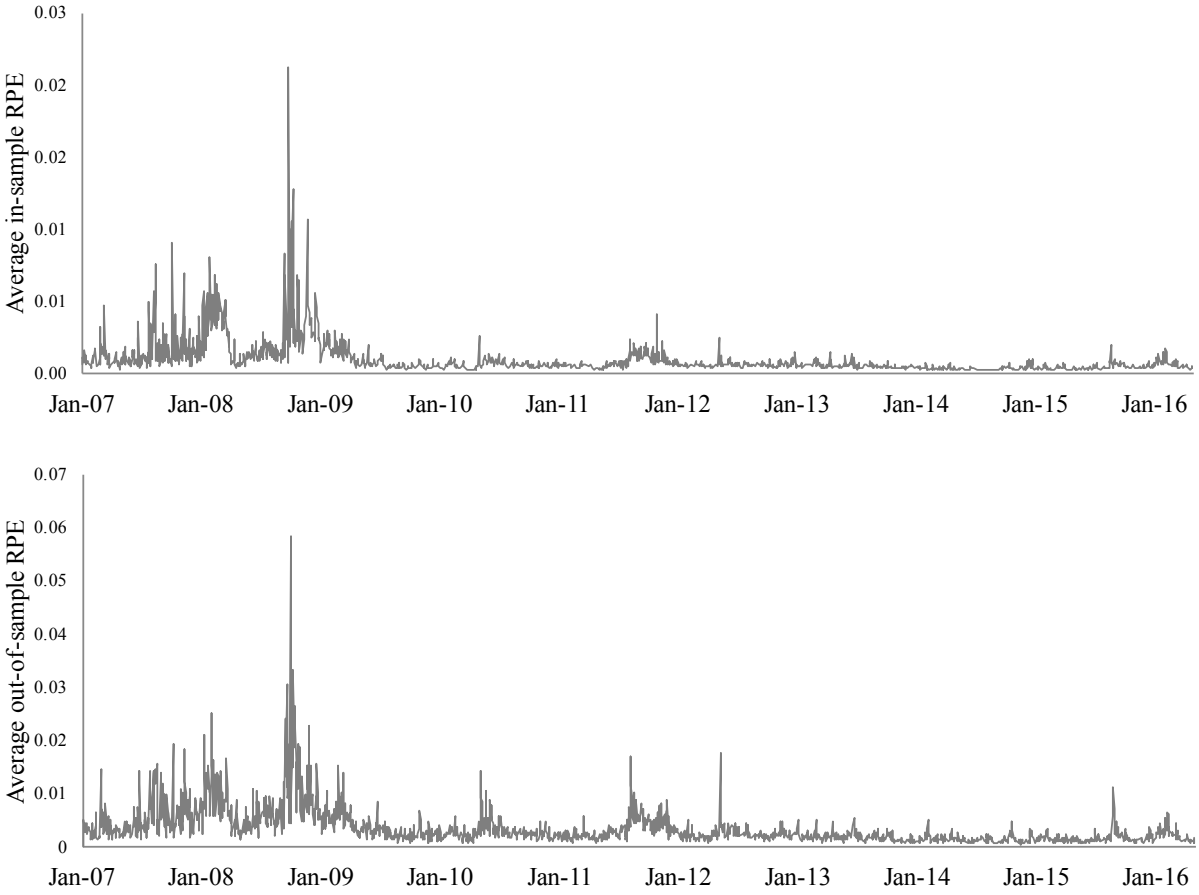

Daily average relative pricing errors of the GH expansion over time (average over all available options and maturities at any given date). The upper graph shows in-sample pricing errors (after removing 17 RNDs with *RPE*s > 2.5%), while the lower graph shows out-of-sample (one trading day forecast) pricing errors, using updated mean and volatility estimates.

Lastly, we observe that the daily pricing performance (measured by *RPEs* across all options on a given trading day) is highly time-varying (*see Figure 1*). Especially during times of crisis, the expansion cannot be sufficiently fit to option data and RPEs are relatively high. This again can relate to the argument about mean and volatility estimates, as provided above.

4.3 Out-of-sample pricing performance

A first step and an intuitive method of testing whether RNDs are stable over time is to check for their out of sample pricing performance. If errors are significantly higher out-of-sample than in-sample, while controlling for the level of the underlying, mean, volatility and other inputs (generally only keeping the expansion coefficients constant), we can conclude that the RNDs

are not stable. Necula, Drimus and Farkas (2016) indicate good out-of-sample performance (by linear extrapolation of parameters from the previous day) compared with other methods, but errors on average are still much higher than in-sample. For our purposes, we differentiate between three methods: first, we obtain out-of-sample option price estimates by using previous-trading-day mean, volatility and GH-expansion coefficients to price options 'today', whereby we must keep in mind that we use the RND estimates corresponding to the correct maturity. Second, we use 'todays' volatility estimates with previous-day mean and GH expansion coefficients. Third, we update both volatility and mean estimates and use previous-day GH expansion coefficients. In all three methods, we use updated levels of the underlying as well as updated dividend yields and risk-free rates.

Table 3 shows that our out-of-sample pricing performance is significantly worse than in-sample performance. We thus conclude that RNDs are not stable enough to ensure a reasonable performance of the model over time, given constant GH expansion coefficients. This is interesting as it points to changing shapes of the GH-RNDs and is exactly the point which we will study later. (We recognize that the time differential between two adjacent trading-days is a large time span considering today's market speed. It would obviously be interesting to study if results would differ when using smaller time intervals, i.e. intraday or tick quotes. As we are constrained to the dataset at hand, we however cannot present an answer to this and leave it to future research.)

Table 3. Out-of-sample pricing errors of the GH model.

	Time to Maturity			
	< 0.25	0.25 - 0.5	0.5 - 1	All
μ, σ not updated	0.56%	0.48%	0.51%	0.53%
only σ updated	0.28%	0.34%	0.41%	0.33%
μ, σ updated	0.24%	0.27%	0.29%	0.26%
in-sample RPE	0.07%	0.06%	0.09%	0.07%

Average relative out-of-sample pricing errors of the GH expansion when using previous-day RNDs for option valuation. Previous day here refers to the previous trading day, i.e. the time difference can exceed 1 day. We differentiate between three methods: first, we use all RND parameters from the previous day. Second, we only use historical GH expansion coefficients and mean but use the current day volatilities. Lastly, we use historical GH expansion coefficients with fully updated mean and volatility estimates. All non-RND-related variables (underlying, risk-free rate, etc.) are updated for pricing. For comparison, we again report in-sample pricing errors.

However, we also observe that while we find an out-of-sample performance comparable to the reference literature using the first method (we do not report *MAE* in the table, but we find the MAE for short maturities to be 0.86, vs. 0.81 in Necula, Drimus and Farkas, 2016), the method of using updated estimates of volatility yields much lower errors. Further including

updated mean estimates decreases pricing errors even more, but only has a much smaller effect. Generally, mean and volatility estimates thus, must to some degree be compatible with GH-coefficients across different days. From a market participant's perspective, this may be imperative when it is necessary to price new options on a day and in a market, where there are not enough options traded such as to sufficiently fit the complete RND (but enough to fit mean and volatility).

Combining above insights, this out-of-sample pricing behaviour indicates that accounting for shocks in the underlying and to implied volatility is still not sufficient to capture changes in the full IVS. Instead, large out-of-sample errors (using method 3) imply that the shape of the RND is highly time-varying and there might be higher-order risk-factors associated with option portfolios. This has implications for risk-management, which will be discussed in the context of the time-series study of shape-measures.

Also, we want to stress that the out-of-sample pricing performance is highly time varying (see *Figure 1*), too. Consequently, the parameter stability of the GH-expansion based RNDs between individual trading days must also be time-varying, and might be linked to market uncertainty and the general level of volatility within the market. We will later study this behaviour in more detail and link it to the volatility of changes in risk-neutral skewness.

4.4 Tail behaviour and additional filtering

We previously saw that for some observations, in-sample pricing errors have been very high. However, in order to interpret the RND estimates in a meaningful way, we have to be confident that they are representative of the 'true' RND expected by the market. Hence, we must apply additional filters to remove observations (i.e. trading day and maturity combinations) where pricing errors are too high (meaning that the GH expansion does not converge to the true RND). Now we cannot simply filter pricing errors based on mean absolute errors, as this would not consider differing general price levels of options with different times to maturity. Instead, we propose to filter observations based on their relative absolute pricing errors and suggest removing RND observations that have an RPE larger than 2.5% of the average option price in that observation. We believe that this should be sufficient to ensure with some confidence that the estimated RND is the one 'expected by the market'. Using a more stringent filtering (i.e. even lower pricing errors as a cut-off) would result in some trading-days with no RND estimates left.

Also, we filter out longer maturities (> 3 months), because shorter ones are the most liquid, as well as RNDs of options with less than a week to maturity.

This method of filtering also has an additional advantage, which is linked to the tail behaviour of Gauss-Hermite RNDs. From *Appendix E*, we see that these densities can have multiple peaks, which we found to be quite common among our estimated RNDs. Generally, those in Necula, Drimus and Farkas (2016) also sometimes exhibit multiple peaks (especially for short maturities) and this kind of behaviour has been observed for Gram-Charlier Type-A kind of expansions in earlier papers (e.g. Jondeau and Rockinger, 2001), too. We believe that in the Gauss-Hermite model, this can result from a) the fact that we truncate the expansion at n=20 which doesn't yield full convergence, b) discreteness in option data, which is especially pronounced on some days and less so on others, and c) the lack of option data in the tails of the distribution.

Now in itself, the property of having multiple peaks is not problematic for an RND, even if these shapes are rather non-standard. Interestingly, the RND in *Appendix E* still yields very low pricing errors and adheres to the characteristics of a probability measure as outlined in *Chapter 3.1* (except for the probability mass, which is only very close to 1). Since we estimate the parameters of the RND by fitting them to the centre of the distributions, and since for pricing only cumulative densities starting at the smallest available strike matter, tail behaviour can be somewhat arbitrary, even after putting restrictions on the densities. Now, while these RNDs produce very low pricing errors, we recognize that 'humps' that are very far in the tail pose problems when estimating skewness and kurtosis. We see this from *Equations (7)* and *(8)*.

This is due to the fact that small absolute values of x (i.e. in the range $-1 < x < 1$) generally do not contribute much to the estimated higher moments, since x^3 and x^4 will naturally become quite small in this interval. Instead, skewness and kurtosis are most heavily influenced by values outside $-1 < x < 1$, for which we however generally do not have many option quotes. In that sense, the GH expansion might not be the best method for studying higher moments, as its tail behaviour is hard to motivate and the sometimes multi-peaked RNDs can prove to be problematic. The method based on general extreme value theory proposed in Figlewski (2010) might thus provide better estimates. A potential solution can be the study of confidence intervals of the distribution instead, which might be easier to interpret and is less prone to errors.

Lastly, note that the above critique mostly applies to kurtosis, as skewness also measures the asymmetry of a distribution and thus can still be interpreted in a meaningful way. Also, the filtering for *RPE*s and for shorter times to maturity and pricing errors as outlined above appears to remove the RNDs with the most extreme tail shapes. We thus proceed by studying if our remaining RNDs and the corresponding estimates of skewness and excess kurtosis can be used

to create a meaningful time series and if they can be used to replicate some findings in the literature.

Table 4. Descriptive statistics of estimated skewness and excess kurtosis.

	Mean	Median	Standard Deviation	Min.	Max.	Count
Skewness	-2.05	-2.06	0.51	-6.75	-0.34	14,043
Excess Kurtosis	8.48	8.12	3.47	0.83	36.94	14,043

Summary statistics of skewness and excess kurtosis estimates that have been obtained from the previously estimated RND's (after filtering as outlined in *Chapter 4.4*) and calculated accordingly to the expressions derived in *Appendix D*.

Filtering as above, we are left with 14,043 observations. Using these, we can directly estimate skewness and kurtosis according to *Equations (9)* and *(10)*. We provide summary statistics on these moments in *Table 4.* We see that skewness is always negative, which implies a fat left tail. Excess kurtosis is consistently above zero, and thus there is always more weight in the tails.

4.5 Time-series indices of higher order moments

We previously obtained skewness and excess kurtosis estimates for multiple option chains with different maturities for each trading day. To create a meaningful time series of moments of the RND however, we must somehow aggregate them. To do so, we have to make sure that the RNDs included in the index are somewhat comparable and well behaved. Thus, the index will only consist of RNDs that passed some strict filters. For the sake of clarity, we will here present these again.

As outlined above, we first filter individual options that have a bid-price of zero or no quoted implied volatility, options that do not satisfy the no-arbitrage condition (see *Section 4.1*), as well as in-the-money options. Second, we do not estimate RNDs for date-maturity-combinations that have less than 23 different exercise prices left after applying the previous filters. Third, to make sure that the estimated RNDs are close to the 'true' RNDs, we filter out those date-maturity-combinations on which the corresponding estimated RNDs yield average relative pricing errors larger than 2.5%. Lastly, we only include RNDs with maturities ranging from 1 week (shorter maturities might introduce pricing problems shortly before expiration) up to 3 months (longer maturities are less liquid and yield more extreme tail-shapes, as already discussed earlier) in the index. Given these filters, we can now combine the remaining RNDs according to the definition given in *Chapter 3.5 (Equation 11)*. Before studying these indices, we however want to highlight some important features of the skewness and kurtosis index here.

First, note that we omit the kurtosis index from here on, since it is highly correlated (85%) with that of the skewness index (see *Appendix F*). There is a clear relationship: whenever the skewness index increases, so does the kurtosis index. This is due to the obvious link by definition of both indices. Thus, kurtosis mostly stems from a fat left tail (since a higher skewness index reflects higher *negative* skewness). Also, as Bakshi, Kapadia and Madan (2003) pointed out, adding kurtosis terms does not significantly improve option pricing models, which is why we believe that kurtosis does not add much additional information in our analysis. The time series created for skewness and that of the S&P 500 index are shown in *Figure 2.*

Lastly, note that our index of skewness follows the CBOE SKEW index rather closely (see *Figure 2*), with both indices exhibiting a correlation of 63% over the whole time-period. The overall trends are similar in both series. Only while our index is more volatile in the earlier part of the series (2007 - 2008), the opposite is true for the latter part of the series (2014 - 2016). From 2010 to 2013, both series are very similar. The differences in the earlier and later parts of the series might stem from the number of options included in our skewness index. During 2007 and 2008, we have quite a low number of RNDs that pass our filters, such that the index on some days only consist of one or two often time-varying maturity bands. Towards the end of our dataset however, we have many more available option quotes and RNDs that pass our filters, such that the index combines many RNDs which makes the index less volatile.

Still again, both CBOE SKEW and our skewness index closely follow each other. This relationship is especially striking, since the CBOE SKEW is calculated in a completely different way with a method based on that outlined in Bakshi, Kapadia and Madan (2003). This suggests that in fact we can use the RNDs estimated from a Gauss-Hermite expansion method to create a meaningful skewness index.

We proceed by studying some dynamics of risk-neutral skewness using this index. What follows is a combination of graphical analysis of skewness over time, which is supplemented by some statistical tests for its time series properties.

4.6 Qualitative analysis

How can we interpret the skewness index? To abstract from the complexity of previous sections, we now start with a simple graphical analysis of the time series of skewness. First however, let's again define what skewness means. As already mentioned above, skewness is a measure of the degree of asymmetry of a distribution about its mean. Importantly, a distribution with a negative (positive) skewness can imply either fatter or longer tails on the left (right) side of the distribution when compared to its right (left) tail. However, a distribution with a skewness

Figure 2. Gauss-Hermite and CBOE skewness indices and the S&P 500 (January 3, 2007 – April 29, 2016).

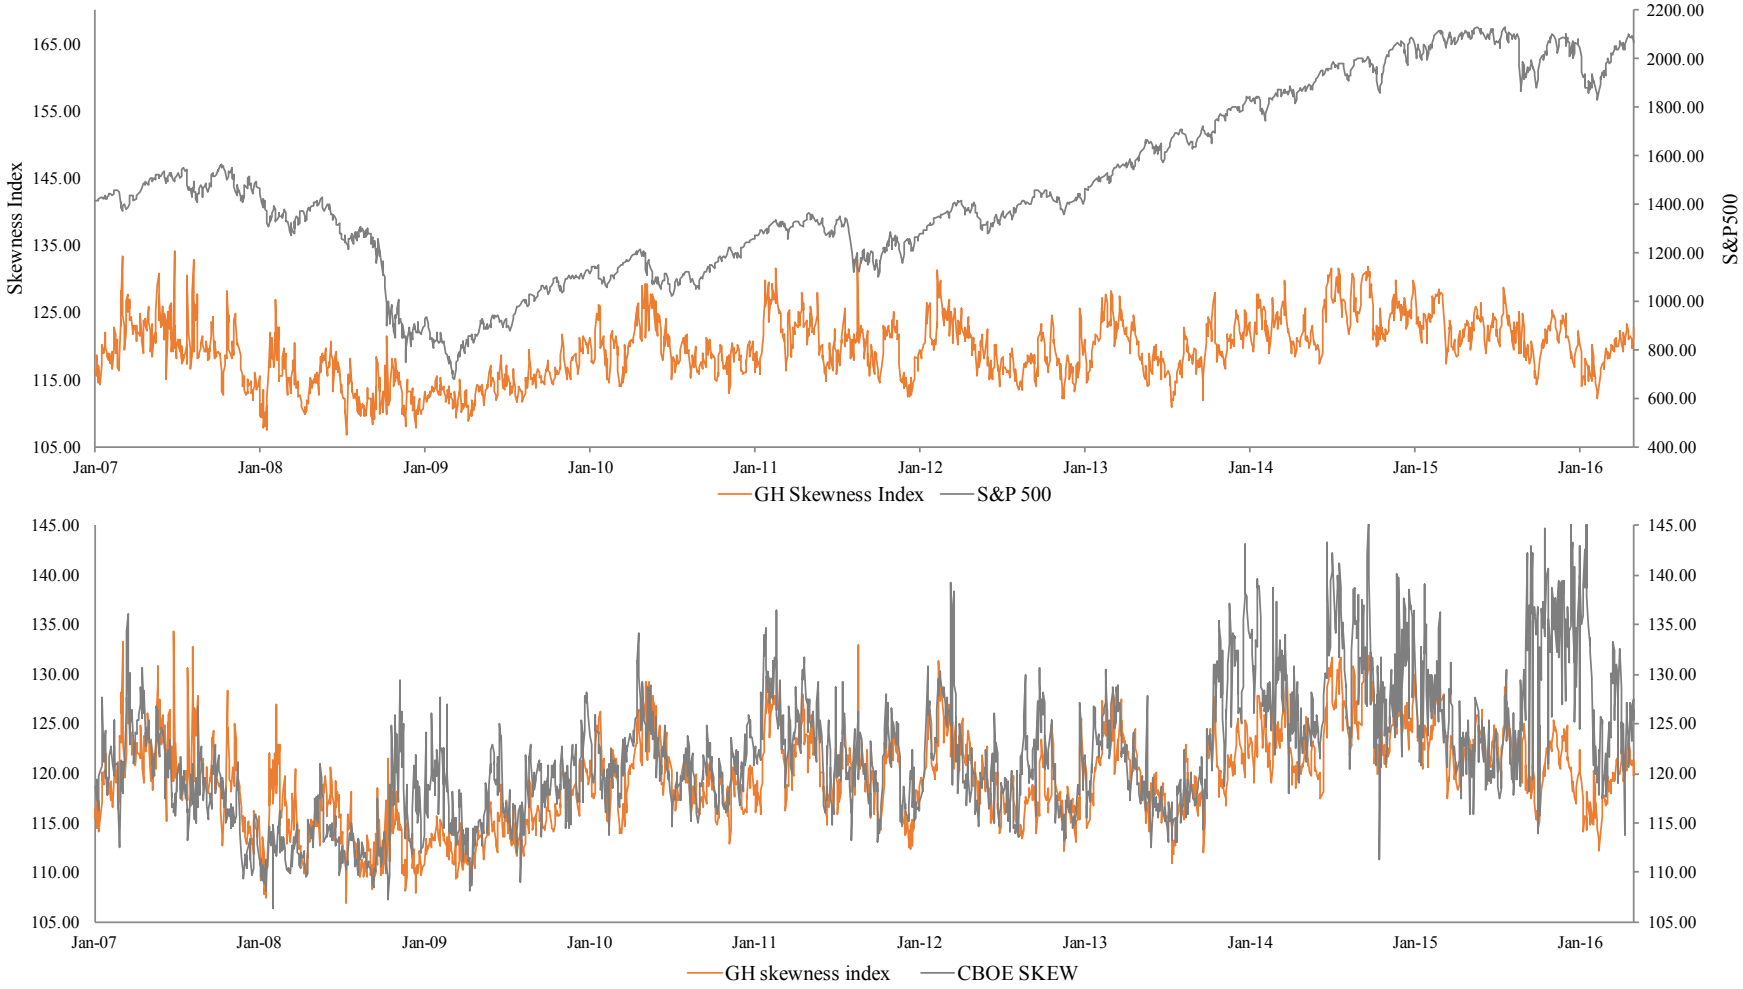

These graphs show the Gauss-Hermite skewness index and the S&P 500 (upper graph) as well as the Gauss-Hermite skewness index and the CBOE SKEW (bottom graph) between January 3, 2007 and April 29, 2016.

of zero may not necessarily be a strict symmetric distribution but may imply that the asymmetric nature of one side of the tail balances out with the other. Thus, a distribution with a fatter left tail and longer right tail may still have a skewness of value zero. Skewness thus essentially measures the relative weights of both tails, compared from the origin (as we defined skewness to be normalized).

Why is skewness important? An interesting calculation about the effect of skewness on the probability of negative log-returns of 2 and 3 standard deviations is shown in CBOE (2010). A higher skewness increases tail risk dramatically. Thus, while implied volatility scales risk-neutral distributions evenly across all strikes, increasingly negative risk-neutral skewness captures market participant's assessment that an extremely negative market event is becoming more likely. We will now study its evolution over time.

Generally, from *Figure 2* we see that the skewness index is a volatile mean-reverting time series. It has a minimum of 106.9 (skewness of -0.69) and maximum of 134.3 (skewness of -3.4), with only very few days with such extremely skewed RNDs. We see that the index is always above 100, which means that the average skewness of options with maturities from one week to 3 months is negative over the whole time-period, suggesting a fatter or longer tail (compared to the Gaussian distribution) on the left side of the risk-neutral densities. It is also especially striking that the level of skewness was rather low during the 07'-08' financial crisis while spikes in skewness seem to be preceded or occur simultaneously to large losses in the underlying. Potentially, this can be explained by the fact that implied volatility was already very high during the financial crisis period and that from *Figure 3* we see that extreme levels of VIX typically go along with only moderate levels in the skewness index. This has already been pointed out in a paper by the Chicago Board Options Exchange (2010) and it suggests that in times of high economic uncertainty the implied asymmetry and fatness of tails of the underlying is rather small, while the whole distribution is stretched simply by a higher level of expected volatility. We proceed to studying this behaviour for individual sub-periods of high volatility and market-relevant events in more detail below.

The U.S. financial crisis (2007 –2009) was a devastating one and exhibited periods of extreme market volatility with many stock-indices dipping to their lowest values in decades, liquidity drying up and implied volatility (VIX) spiking to all-time highs. We again observe that during this period, the skewness index in fact exhibited its minimum values, with a daily average of 113.39, compared to 119.38 of the whole dataset. Now, let's look at individual trading days. On September 15, 2008, the day Lehman Brothers filed for bankruptcy, the S&P 500 lost nearly 5% to its previous day closing price. This event also caused a spike in the

skewness index, which jumped from 110 to 114. In the turbulent following days, the S&P 500 suffered further heavy losses, losing another 5% on September 17 and 8.8% on September 29 (on this day, the U.S. House of Representatives rejected the governments' bailout plan). These events are also reflected in the skewness index, which jumped by 7.8 and 3.6 points, respectively. However, these extraordinary market events do not lead to an abnormal increase in the general level of the skewness index, which was still relatively low compared to other time periods. This relates to the argument about the connection between high volatility and low skewness periods. Interestingly, subsequent losses in the S&P 500 did not induce similar jumps in the skewness. While the S&P 500 declined by 5.9% on October 7 and suffered a 6.3% loss on October 22, the skewness index in fact decreased by 7 and 1.4 points, respectively. Thus, while some adverse events and negative returns in the underlying are met with increased skewness, this is not true of others during the period of the financial crisis (see *Figure 4*).

Figure 3. Skewness index and VIX.

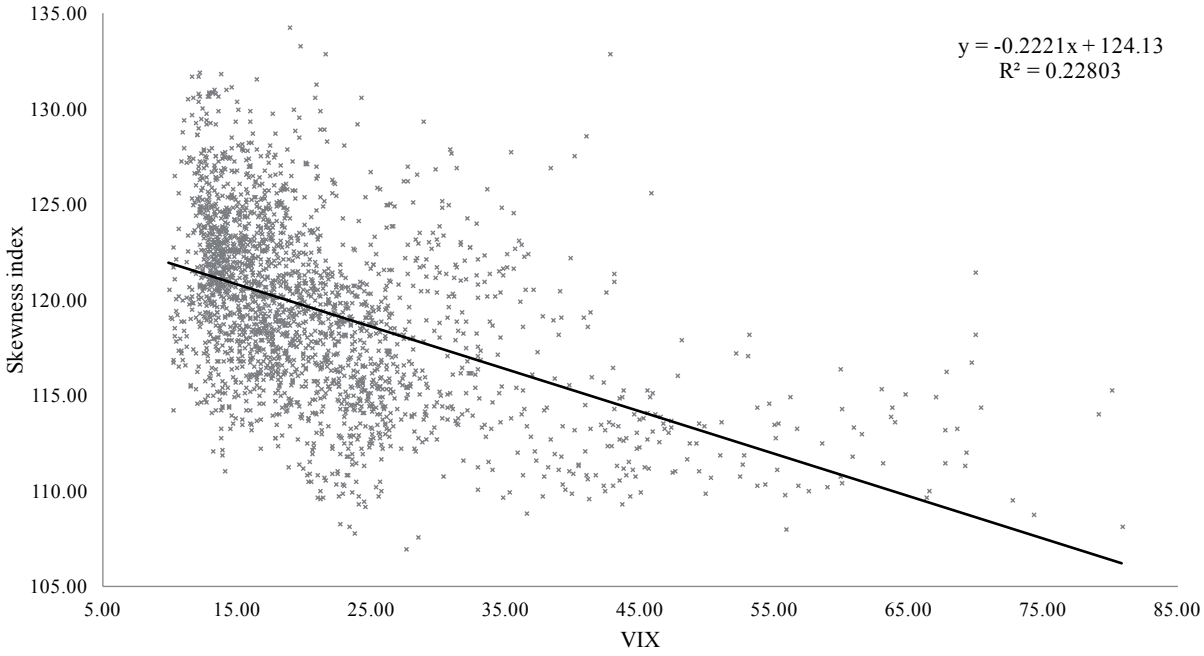

This graph shows daily levels of the skewness index and corresponding levels of the CBOE volatility index (VIX).

After the financial crisis, there of course have been other times of market turmoil which corresponded to highly volatile events within the world economy. First, on May 6, 2010 U.S. equity indices suffered one of their largest intraday losses (in a matter of minutes) in what has become widely known as the Flash Crash. While the S&P 500 at its close had lost 3.3%, the skewness in fact decreased on that day and only jumped to a slightly higher level the day after. Second, a similar pattern emerged after August 8, 2011 (Black Monday), the day the U.S. sovereign debt rating was downgraded from AAA to AA+ by Standard and Poor's and the S&P

500 suffered consecutive large losses for several days. While the S&P 500 lost nearly 7% on that day, the skewness in fact did decrease and was rather flat the following days. Only following the 5% market drop on August 18 and a worsened economic outlook did the skewness spike by more than 6 points, but it immediately dropped to its previous levels the day after.

It seems that adverse market events do not always correspond to spikes in risk-neutral skewness. In fact, skewness is sometimes persistently low (<110) for several trading days during the financial crisis. This for example was the case between January 3 and January 17, 2008, which corresponded to the onset of the crisis when investors were mulling over a looming recession or in November 2008 when federal bailouts for U.S. auto-makers were rejected by the republicans.

From the above analysis, we conclude that there is no clear relationship between extreme returns of the underlying and swings in the level of risk-neutral skewness, which in itself is very volatile. We also see that the average skewness of options with maturities from one week to 3 months is consistently negative over the whole time-period. Furthermore, we find that the level of the skewness was relatively low during the 07'-08' financial crisis, which however corresponds to the fact that extremely high implied volatilities usually correspond to less asymmetric RNDs. We will study a few of these observations more closely in the following section.

Figure 4. Skewness index and S&P 500 (August – December 2008).

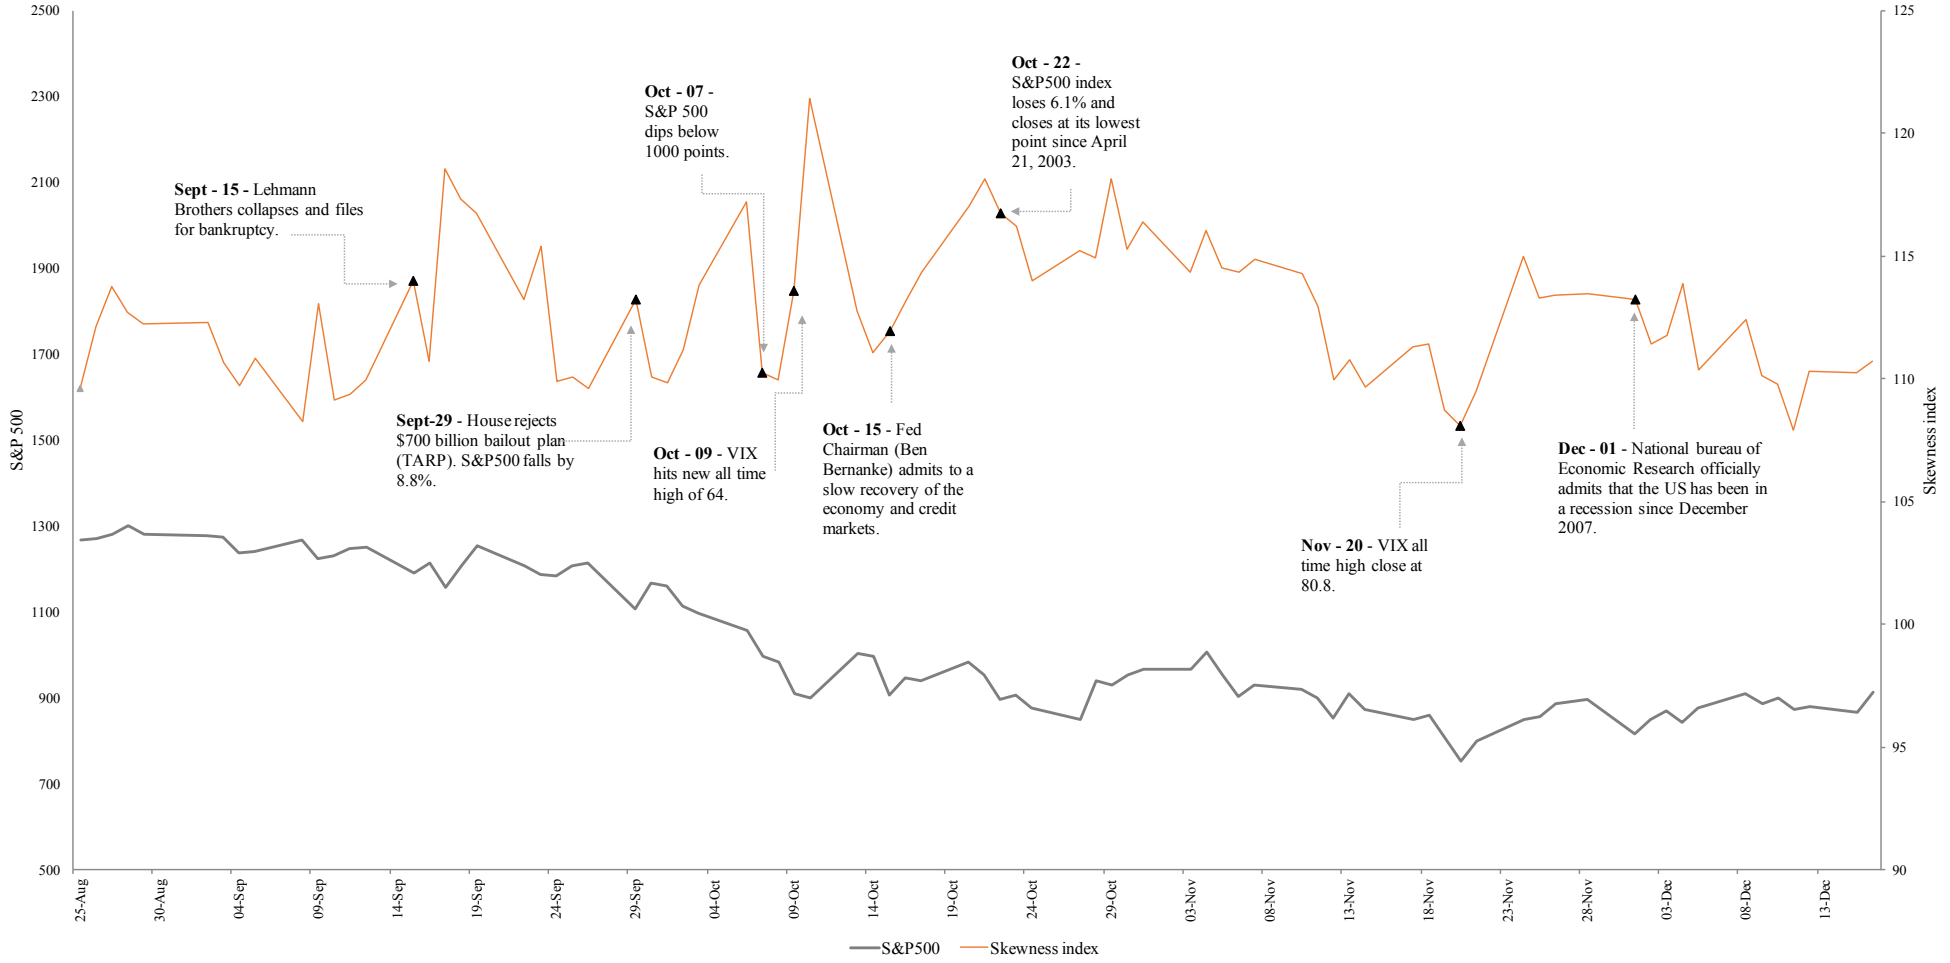

This graph shows the Gauss-Hermite skewness index and S&P 500 during the most volatile period of the financial crisis (August 2008 – December 2008) and highlights market-relevant events.

4.7 Time series statistics

In the following, we analyse the time series behaviour of the skewness index in a more statistical framework. We also ask if there is a significant relationship of changes in the underlying and that in risk-neutral skewness.

A first step is to check for autocorrelation in the time-series of the skewness index. From *Figure 5,* we see that there is significant autocorrelation present in the series, and it is decreasing slowly after many lags. This means that high levels of skewness typically are followed by high levels of skewness, and low levels are associated with previous low levels. Similarly, the same is true for the VIX. A natural question to ask is where the variation in skewness and VIX stems from. For the VIX, a well-documented behaviour is the so-called *leverage effect*. It postulates that implied volatility increases for lower levels of the underlying (for example Geske, 1979). Thus, some variation in the VIX series can be attributed to returns in the underlying. We show this in the first regression in *Table 5,* where we denote the S&P 500's log return by $\Delta S_t = \ln(S_t/S_{t-1})$. Regressing today's level of VIX on the one day lagged VIX and today's log-return in the underlying, we confirm that both coefficients are highly significant and the high associated R^2 implies that a large part of the variation in the VIX can be explained by these coefficients. Thus, negative returns in the underlying indeed typically go along with increases in implied volatility. A similar picture emerges for the skewness index when regressing it on the associated S&P 500 return and the previous day's skewness (see the second regression in *Table 5*). However, the R^2 of the equation (83%) is not as high as that of the equation for VIX. Thus, a larger part of the variation in the level of skewness cannot be explained by these coefficients.

While the above method is the standard way of checking for the leverage effect, we believe it would be more interesting to remove the autocorrelation in VIX and skewness and see how much of the remaining variation can be explained by returns in the underlying. To remove autocorrelation, we take the first difference in the skewness series, i.e. $\Delta\bar{\gamma}_{1,t} = \bar{\gamma}_{1,t} - \bar{\gamma}_{1,t-1}$, and do the same with the VIX series. Both series still exhibit some significant autocorrelation at early lags, however autocorrelation dies out quickly (see *Figure 5*). For the sake of simplicity and readability, we do not study the autocorrelation further, as its effect is negligible (see *Table 5*, specification 5). Thus, we regress first differences in VIX and skewness solely on log-returns of the underlying in equations 3 and 4 in *Table 5*. Here, the picture becomes much clearer. While the coefficient assigned to log-returns of the S&P 500 is highly significant in both equations, the underlying's returns explain 72% of the variation in the first difference of the VIX series, but 0% of that of skewness. Hence, daily changes in risk-neutral

skewness cannot be explained with corresponding returns of the underlying. We also show this in *Figure 6a*, from which we can see that there is no clear relationship in the tails, either. Thus, not even extreme movements in the underlying affect the skewness index in a consistent fashion. This is in line with our previous graphical analysis.

Figure 5 - Autocorrelation plots of skewness index and VIX.

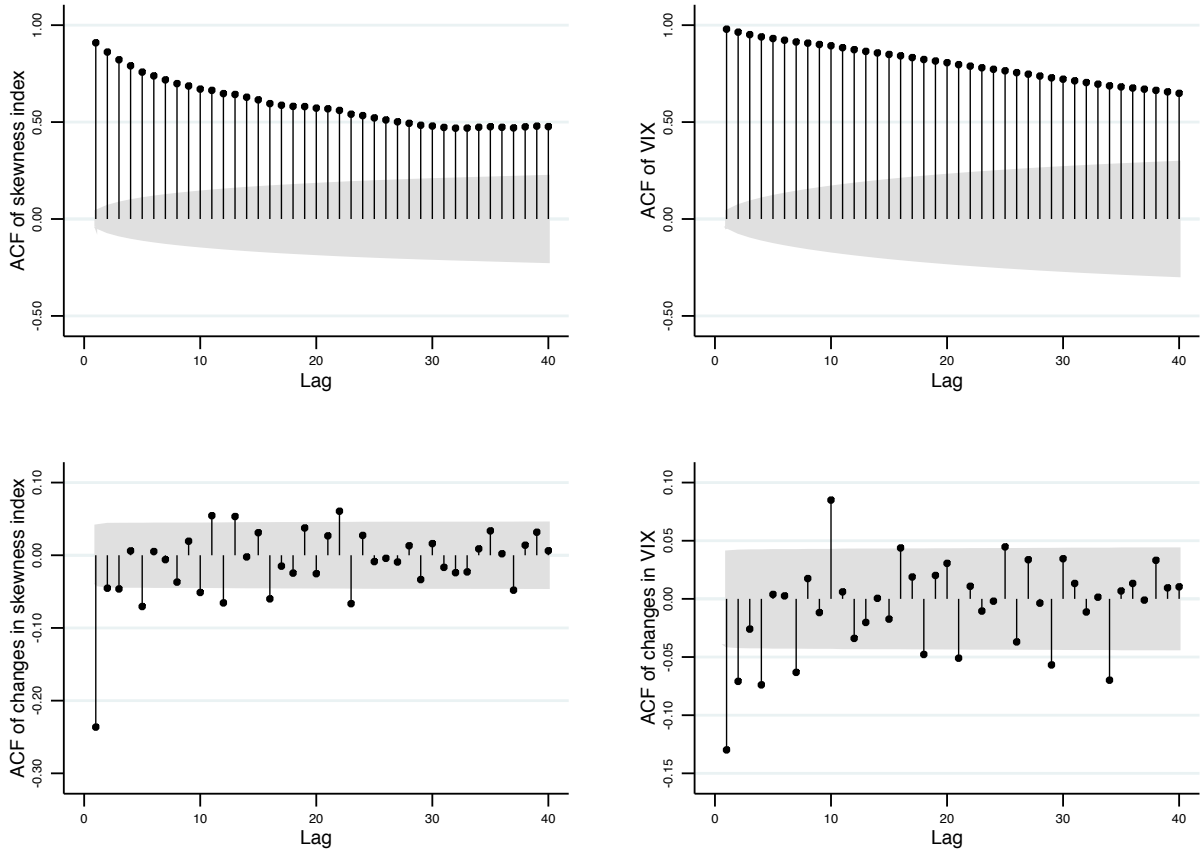

These graphs show the autocorrelation functions (ACF) of the skewness index (upper left), VIX (upper right), first differences in the skewness index (bottom left) and first differences in VIX (bottom right) for up to 40 lags.

To fully reject an effect between movement in the underlying and the skewness, we must also study for lagged effects between both variables. For this purpose, we propose to study these relationships in the framework of a vector autoregressive (VAR) model. Such a model fits changes in both underlying and the skewness index on *previous* (lagged) changes in both variables. To study for such a relationship, we must specify the maximum number of lags included in that model. This is tricky, because one must consider both statistical information criteria and economic considerations. Often, there is a trade-off between the models' performance and its complexity when it comes to the numbers of lags included. To make things worse, we find that different information criteria suggest different optimal lag-lengths for a model containing both series of changes in skewness and the underlying (see *Appendix G* for

Table 5. Linear regressions[1].

No.	Dependent Variable	Intercept	$\sigma_{t\text{-}1,VIX}$	$\bar{\gamma}_{1,t\text{-}1}$	ΔS_t	$\lvert\Delta\sigma_{t,VIX}\rvert$	$\lvert\Delta\bar{\gamma}_{1,t}\rvert$	$\Delta\bar{\gamma}_{1,t\text{-}1}$	$\sigma_{\gamma,t,EWMA}$	R^2
1	$\sigma_{t,VIX}$	0.35*	0.98*		-125.1*					0.99
2	$\bar{\gamma}_{1,t}$	10.86*		0.91*	7.81*					0.83
3	$\Delta\sigma_{t,VIX}$	0.02			-125.42*					0.72
4	$\Delta\bar{\gamma}_{1,t}$	0.001			7.705*					0.00
5	$\Delta\bar{\gamma}_{1,t}$	0.003						-0.24*		0.06
6	$\overline{RPE_{t,1}}$	0.001*				0.01*				0.70
7	$\overline{RPE_{t,3}}$	0.002*					0.001*			0.10
8	$\overline{RPE_{t,3}}$	-0.001*							0.315*	0.27

This table shows the results of several regressions used to test dependencies involving the skewness index. The dependent variables are reported in the left column. $\bar{\gamma}_{1,t}$ denotes the level of the skewness index, $\Delta\bar{\gamma}_{1,t}$ is the first difference in that series, ΔS_t is the log-change in the S&P 500, $\sigma_{t,VIX}$ denotes the level of VIX, $\Delta\sigma_{t,VIX}$ is the first difference in that series, and $\sigma_{\gamma,t,EWMA}$ denotes the estimated EWMA volatility of the skewness index. $\overline{RPE_{t,1}}$ and $\overline{RPE_{t,3}}$ are the daily average relative pricing errors of the GH-expansion, calculated with previous day as well as updated mean and volatility estimates, respectively. (*) denotes significance at the 1% level. Empty cells mean that a variable has been omitted in the corresponding equation.

Figures 6a and 6b. Correlation of changes in the skewness index with S&P 500 returns (left) and changes predicted by an AR(1) process (right).

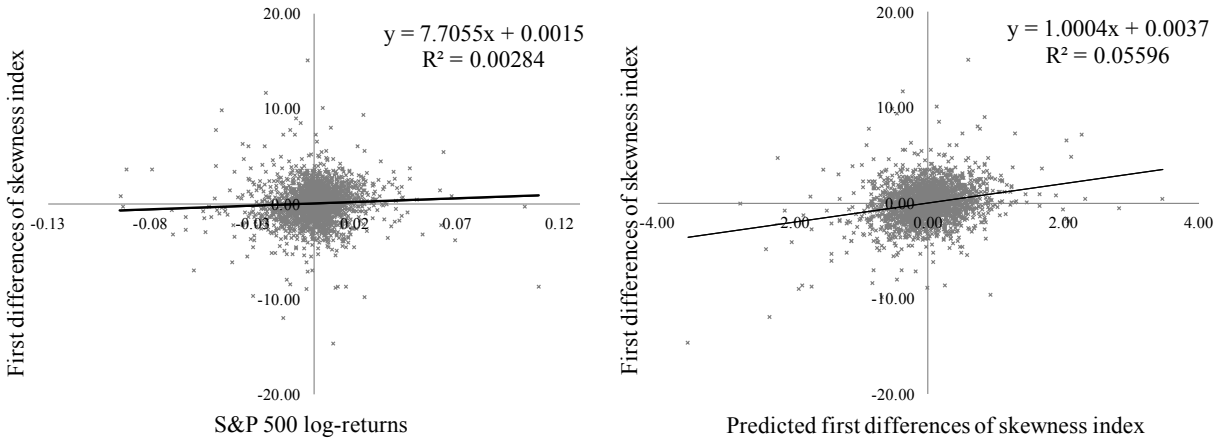

These graphs show the correlation of first differences in the skewness index against log-returns in the S&P 500 (left graph), and the correlation of first differences of the skewness index predicted by an AR(1) model (see *Table 5*, equation 5) and actual changes in the skewness index (right graph).

Stata outputs). While the Akaike information criterion (AIC) suggests the inclusion of up to 18 lags, the Bayesian information criterion (BIC) suggests only 2 lags (it does penalize the

[1] All regressions in Table 5 are (multiple) linear regressions of the form $y = X\beta + u$, where y contains observations of the dependent variable, X contains observations of the independent variables and the intercept, β is the vector of coefficients and u contains the stochastic error terms.

inclusion of additional parameters more than AIC does), and the Hannan–Quinn information criterion (HQIC) suggests 5 lags. We suggest sticking with 5 lags, which balances both AIC and BIC, is easy to interpret (5 trading days means one week) and should be more than enough time for all information to diffuse between markets. Now, in the context of a VAR model we can also test the hypothesis, that all coefficients of all lags (that are included in the model) of variable x on variable y are zero. This is equivalent to the null hypothesis of no Granger causality of x on y. Checking for Granger causality in this fashion, we find that we cannot reject the null hypothesis of no Granger causality of lagged changes in the underlying on risk-neutral skewness (p-value of 0.17). Thus, changes in the skewness-index are not Granger caused by log-returns of the underlying. On the other hand, we can reject the null hypothesis for changes in skewness on changes in the underlying (p-value of 0.013). However, when tracing this back to the VAR estimation results, this is probably due to multiple coefficients of skewness on the S&P 500 that are close to being statistically significant. However, the R^2 of that equation is extremely low (0.02), which is why we also disregard potential economic effects of changes in risk-neutral skewness on returns of the underlying. Also note, that various coefficients of lagged changes in skewness on 'current' changes in skewness are significant. Thus, there might be autocorrelation present in that series. However, we already disregarded the effect to be economically insignificant, and the autoregressive AR(1) model estimated from the data does not appear to predict movements in the skewness very well, as is shown in *Figure 6b*. Including more lags does not improve results.

4.8 Skewness and pricing errors

The fact that skewness cannot be explained by returns of the underlying points to the interpretation that skewness may constitute another risk-factor when it comes to option-pricing or risk-management of option portfolios.

We already saw earlier that the pricing errors of the GH-expansion can be significantly reduced when accounting for updated volatility estimates (see *Chapter 4.3*). We also show this in regression 6 of *Table 5*, where we regress the daily average relative pricing errors (out-of-sample, using previous-day's mean and volatility) across all RNDs on absolute changes in the VIX. The associated coefficient is highly significant and positive, and absolute changes in the overall level of VIX explain 70% of the variance of the out-of-sample pricing errors. Thus, whenever VIX jumps, the previous day's RNDs yield high pricing errors when used to price options out-of-sample. Unsurprisingly, VIX is thus a risk-factor for option portfolios.

Now, if one accounts for updated mean and volatility estimates and denotes the resulting daily out-of-sample pricing errors by $\overline{RPE_{t,3}}$, we can check if this remaining pricing risk is to some degree explained by absolute changes in risk-neutral skewness. From *Table 5* (regression 7), we see that the associated coefficient is again both positive and highly significant. However, changes to risk-neutral skewness explain only 10% in the remaining variation of out-of-sample pricing errors (after accounting for updated mean and volatility). Thus, skewness might indeed be another risk-factor for option portfolios, its effect however is relatively small. This is to some degree surprising; however, we also recognize that changes in skewness cannot capture all details of changes in the shape of RNDs.

While the immediate correlation between pricing-errors and the magnitude of changes in the skewness index is limited, we also compare realized volatilities of the skewness index to GH pricing errors. For this we calculate daily estimates of realized volatility of the skewness index using an EWMA[1] approach. To keep things simple, we set the weighting factor to $\lambda = 0.94$, which is suggested by RiskMetrics for daily returns. Let's denote the realized skewness volatility by $\sigma_{\gamma,t,EWMA}$ and check for a connection to $\overline{RPE_{t,3}}$. From Table 5 (regression 8), we see that the out-of-sample pricing errors are significantly positively linked to higher realized volatilities in the skewness index, with the associated R^2 being much higher when compared to the regression including absolute daily changes in the skewness index. Thus, while changes in skewness cannot explain a large portion of remaining pricing errors, GH pricing errors are still usually higher in times of highly volatile skewness.

Lastly, we check for a possible relationship between implied and realized volatilities of the underlying and skewness. As already stated above, estimated skewness is rather small for days with extremely high implied volatilities. For days with low levels of VIX however, the corresponding level of risk-neutral asymmetry is not as clear. While the level of skewness is more spread out for smaller VIX levels, they are still associated with the highest levels of risk-neutral skewness. Interestingly, the same relationship also holds for daily historical volatility estimates of the underlying. We can also estimate these using an EWMA approach (again with $\lambda = 0.94$). It turns out that VIX and realized volatilities are highly correlated (94.5%). Thus, in times of extreme market uncertainty (both historical and forward looking), RNDs are on average more symmetrical than over the whole dataset.

We also find that the skewness index is more volatile in times of high volatility in the underlying. We see this when comparing the EWMA-volatility estimates of the skewness index

[1] Exponentially weighted moving average. We define the daily EWMA volatility estimate as $\sigma^2_{t,EMWA} = \lambda\sigma^2_{t-1,EMWA} + (1-\lambda)R^2_t$, with R^2_t being the the squared log-change in the underlying series at time t.

with EWMA-volatility estimates of the S&P 500. While both series are not highly correlated, we find that volatility in skewness still tends to jump when volatility in the underlying spikes (see *Figure 7*). Thus, while the overall level of risk-neutral skewness on average is relatively low when market volatility is high, the uncertainty in the option markets about the skewness of the future RND rises in those times, too.

Figure 7. Skewness and S&P 500 EWMA volatilities.

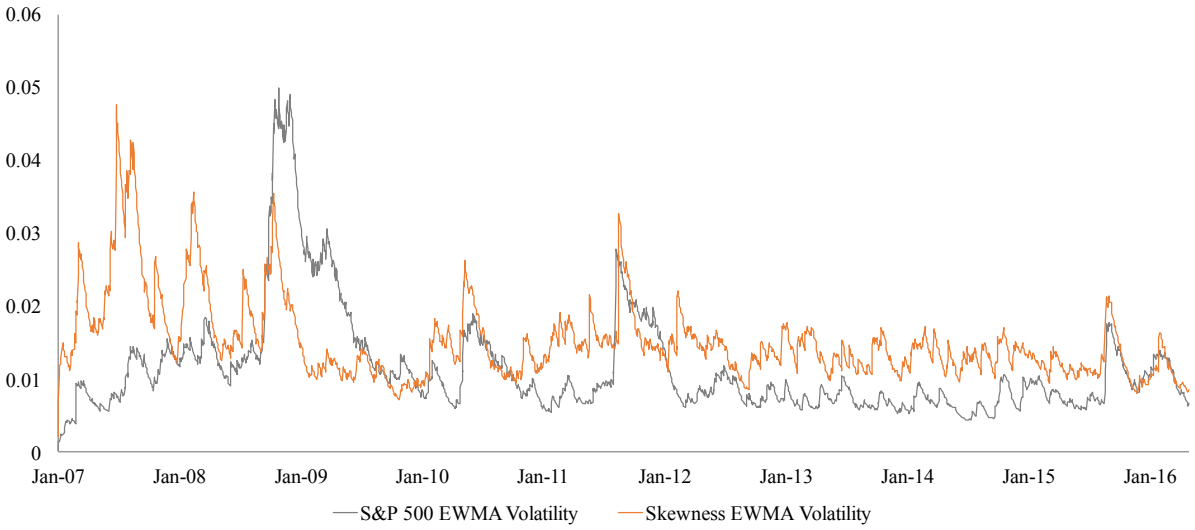

This graph shows daily estimated S&P 500 EWMA volatilities and EWMA volatilities of the skewness index.

5. Discussion

5.1 Implications

We previously studied time-series dynamics of the risk-neutral skewness index extracted from S&P 500 options. We find that skewness is consistently negative, which implies a fat left tail of the RND when compared to the Gaussian distribution. Strikingly, the tail is comparatively thin during the financial crisis and other periods of market uncertainty. We link this behaviour to high implied and realized volatilities during that time, a property that we show to go along with rather low skewness in option prices. We also find that our skewness index is close to the CBOE SKEW index, that it is a mean-reverting series and it exhibits high autocorrelation. However, its daily variability is relatively high and hard to capture in an economic model. We observe that while changes in the skewness index are statistically significantly auto-correlated, the predictive power of that autocorrelation for one-day ahead forecasts is low. Also, returns in the underlying (instantaneous as well as at multiple lags) do not co-move with skewness. We further find similar results for correlations of the skewness index with VIX, even though high volatility states in the market typically go with relatively symmetric RNDs. For 'normal' levels of VIX however, there is no distinct relationship between both series. Moreover, there is no clear and consistent reaction of skewness to market relevant events. We only found that volatility in the skewness series usually increases when realized (and implied) volatility in the underlying does. Combining above insights, it is highly unclear where the variation in the skewness series stems from. This indicates a stochastic process or another yet unknown variable affecting risk-neutral skewness.

We also show that accounting for updated volatility estimates can significantly reduce out-of-sample pricing errors. For short-term options (<3 months to maturity), 50% of pricing errors can be attributed to daily changes in the overall level of implied volatility. Thus, risk-neutral implied volatility is an important risk-factor for options. However, even when employing updated mean and volatility estimates, the remaining out-of-sample pricing performance using previous-day RND shape parameters is still significantly worse than in-sample. This points to time-varying shapes of the estimated RNDs. We link this to variability in the skewness series, which we show to account for some part (10%) of the variability in remaining out-of-sample errors. However, a large part of out-of-sample errors still remains unexplained.

A stochastic risk-neutral skewness that to some degree affects pricing performance is problematic. Thus, our evidence points to skewness as an additional risk-factor for option portfolios, since not all pricing errors can be accounted for with movements in the underlying

(Delta risk) and parallel shifts in the implied volatility surface (Vega risk). These risk-factors are typically employed for risk-management purposes of option-portfolios, but only account for about 50 % of out-of-sample pricing errors in our paper. While the effect of skewness-risk alone is not negligible but also might not be extremely important for risk-management purposes, the unknown source of the remaining variability in the RND after accounting for effects of implied volatilities (and skewness) is even more so. (Since kurtosis has been shown to be highly correlated with skewness, we believe it would not add anything in explanatory power here.) Thus, as Bedendo and Hodges (2009) pointed out, risk-neutral skewness might thus introduce another level to Vega risk.

We further found these results and interpretations to be reflected in previous research on IVS dynamics. An important study by Cont and Da Fonseca (2002) decomposed these dynamics into different risk-factors corresponding to the level, slope, and curvature of the IVS. These correspond to volatility, skewness and kurtosis in risk-neutral space. Their main findings are as follows. First, the IVS is highly dynamic. This is reflected in our assessment of RNDs, since these produce higher out-of-sample errors than in-sample errors. Second, they find that the largest part of IVS variability can be explained by the overall level of volatility. We also find that accounting for this can reduce pricing errors significantly. On a related note, we can confirm the assessment that level effects (VIX) are highly negatively correlated with returns in the underlying. Third, skewness has been shown to be highly auto-correlated, mean reverting, and not correlated to returns of the underlying. Also, it has been shown to have a small (about 3% of the variation) but significant effect on IVS dynamics. Thus, while using a different dataset and a very different methodology, our evidence points in the same directions.

Similarly, previous research on the topic of risk-neutral skewness pointed out links to investor sentiment (see *Chapter 2.3)*. From our analysis, we do not find strong fears or anticipation of market crashes. Risk-neutral skewness was comparatively low in most of those cases, for example right before the Lehman bankruptcy. This relates to findings in Bates (1991).

While we can confirm some findings in the literature, there are also some discrepancies. For example, Gemmill and Saflekos (2000) and Birru and Figlewski (2012) find that RNDs reacted to movements in the underlying, and left-skewness increased only during and after crashes. Contrary to that assessment, we do not find any clear immediate (or lagged) link between returns in the S&P 500 and the level of or changes in skewness. On the other hand, to study possible lead-lag effects of returns of the underlying on implied risk-neutral moments, the use of intraday trading data would most likely draw a clearer picture, since we believe it to

be unrealistic that information would diffuse into option prices (or back) with a lag of a day or more.

Furthermore, since Andreou, Kagkadis and Philip (2012) pointed out that 'rational' sentiment based on economic fundamentals seems to drive option implied skewness, we find it especially striking that skewness in the crisis period was relatively low, while economic fundamentals were deteriorating and the macro-outlook was worsening quickly. We can only hypothesise that uncertainty was on such extreme levels that option prices across all strikes (and not only the left tail) had been associated with high premia during the crisis. Also, from the paradigm of market efficiency, one would expect returns to be independently and identically distributed. A related potential explanation is pointed out in CBOE (2010): markets might simply not view a repeated crash as likely. This also relates to demand-driven option-pricing models, which predict that if markets do not view consecutive extreme losses as likely to happen, then the demand for downward protection and thus for far out of the money put options will be lower. This is reflected in relatively lower premia associated with such options and hence, the implied volatility smile is not as steep on the left side (Bollen and Whaley, 2004). However, these explanations somewhat contradict the notion of volatility clustering.

Lastly, Gemmill and Saflekos (2000) put forward the argument that since estimates of historical return skewness require a large dataset, one instead might possibly use densities implied in option prices for risk management purposes (to forecast future return densities of the underlying). However, our time varying assessment of tail-risk and time-varying out-of-sample pricing performance suggests that the shape of the option-implied RND heavily changes over time. This points to problems when using RNDs in risk-management, since it is not clear which RND one should use. However, we do have to keep in mind that these densities are risk-neutral and not real world densities, thus one must transform them first, as time-varying risk-aversion might in fact contribute heavily to the variation in RND shapes. Also, with our simple analysis, we can obviously not answer the question whether RNDs predict future realized return distributions, which would be required for them to be applicable to risk-management. However, the results in the literature about the usefulness of RNDs for real-world probability forecasting are mixed at best (see *Chapter 2.3*). Thus, risk-neutral densities might not be directly applicable for risk-management purposes, especially as they are highly volatile in times of high volatility in the underlying.

Our results combined with previous research have several implications, as we outlined above. However, they do not draw a fully conclusive picture and several topics are still to be explored. We point to some of them in the following.

5.2 Research outlook

First, due to more stringent legislation and lessons learned from past crises, financial risk-management is becoming increasingly important. It is concerned with the risk entailed in financial portfolios. Obviously, risk management heavily relies on assumptions about the future distribution of the portfolio, its underlying risk-factors and their interdependencies. As discussed above, skewness constitutes an additional, often overlooked risk factor for option portfolios. Also, a significantly large portion of pricing errors cannot be explained by the first three moments of RNDs. Thus, more research would be needed to identify ways of accounting for these kinds of risks.

Second, previous papers have made some progress on decomposing skewness and kurtosis risk in the IVS (see *Chapter 2.2*), which is why we believe that some of these techniques (e.g. Kalman filters) can also be employed to study for dynamics in the GH *coefficients,* which fully parametrize the implied volatility smile. That might draw a clearer picture about the link between skewness and pricing errors.

Third, to extend on the methodology used in our paper as mentioned in the previous point, we also believe that employing intraday option data might facilitate some new insights. This also applies to the study of immediate links between returns in the underlying and risk-neutral skewness, which was pointed out in the literature, but which we however have not been able to confirm with our methodology.

Lastly, while GH-RNDs might not be able to predict future return distributions, the underlying Gauss-Hermite quadrature might still be helpful in risk-management. First, the expansion coefficients can be fit to historical return distributions. This can be helpful to easily calculate Value-at-Risk and other measures analytically (e.g. Pimbly, 2006). Second, some progress has already been made to model multivariate distributions using Gram-Charlier densities (see Del Brio, Niguez and Perote, 2009), and it would be interesting to see if the GH expansion also can be used in such a way, and how it would perform in practice.

6. Conclusion

We use a new method of extracting risk-neutral skewness and a comparatively recent dataset (2007-2016) of S&P 500 index options. We show that while the tail behaviour of RNDs estimated using a Gauss-Hermite approach are sometimes problematic due to the general lack of option quotes far into the tail, stringent filtering (for pricing errors, time to maturity and the number of available strike prices) can still ensure realistic RND estimates.

We derive easy to use formulae for calculating risk-neutral skewness and kurtosis directly from the estimated GH expansion coefficients. Obviously, this is even easier for Edgeworth and Gram-Charlier Type-A based pricing models, which however have sub-par pricing performance, potentially making the GH approach more likely to be used in practice.

We continue by using our estimates to create a time series of option implied skewness (and kurtosis). We find the skewness index to be closely correlated to the CBOE SKEW index, which is computed using a completely different method. We are thus confident that our skewness index can be interpreted in a meaningful way. However, this also means that the GH expansion does not reveal additional information about risk-neutral skewness. We further find that kurtosis is highly correlated to risk-neutral skewness, which not only means that kurtosis mostly stems from a more pronounced left tail (since skewness is usually negative), but also that kurtosis does not add much additional information. We thus solely focus on the time series properties of our skewness index.

We find that the level of skewness implied in option prices is mean-reverting and highly auto-correlated. However, the variation in skewness can neither be attributed to past changes in skewness, nor is it related to returns of the underlying or changes in VIX. Furthermore, we link times of high volatility in skewness to increased out-of-sample pricing errors of the GH expansion. This points to skewness possibly being an additional, stochastic risk factor for option pricing. However, a large part of out-of-sample pricing errors remains unexplained. We can thus replicate some findings in the literature that decomposed dynamics in the implied volatility surface into different risk components (level, slope and curvature of the IVS). Nonetheless, our paper is only a very small first step in this direction and lacks the detail and sophistication of most methods employed in previous papers. We also believe that a more detailed dataset including intraday quotes would improve the assessment of some relationships in the data. Still, our findings in combination with previous literature have obvious implications for risk-management.

Lastly, since options are a way of trading risk, implied risk-neutral densities can provide policymakers with information about aggregate market expectations of future return

distributions and have been linked to market sentiment. Thus, risk-neutral skewness can be used by central banks to assess sentiment and the perceived credibility of monetary policy, which is important for monetary forward-guidance. Similarly, implied RNDs (and therefore their implied skewness) can be used by market participants who wish to base their trading activities on divergences between their own and market expectations. However, to better interpret such measures at any given time, a prior study of their dynamics over time is necessary. While we only pursue this for S&P 500 options, the GH expansion and our measures of skewness and kurtosis could obviously also be used for options with other underlyings.

7. References

Alonso, Francisco, Roberto Blanco, and Gonzalo Rubio, 2005, Testing the Forecasting Performance of IBEX35 option-implied risk-neutral densities, Documentos de Trabajo, No. 0504, Banco de España.

Andersen, Torben G., Nicola Fusari, and Viktor Todorov, 2015, Parametric Inference and Dynamic State Recovery from Option Panels, Econometrica 83 (3), 1081-1145.

Andreou, Panayiotis C., Anastasios Kagkadis, and Dennis Philip, 2012, Investor Sentiments, Rational Beliefs and Option Prices, working paper, Durham University Business School.

Bachelier, Louis, 1900, Théorie de la speculation, Annales scientifiques de l'École Normale Supérieure, Sér. 3 (17), 21-86.

Bakshi, Gurdip, Nikunj Kapadia, and Dilip Madan, 2003, Stock Return Characteristics, Skew Laws, and the Differential Pricing of Individual Equity Options, Review of Financial Studies 16 (1), 101-143.

Bates, David S., 1991, The Crash of '87: Was it Expected? The Evidence from Options Markets, The Journal of Finance 46 (3), 1009-1044.

Bates, David S., 2000, Post-'87 Crash Fears in S&P 500 Futures Options, Journal of Econometrics 94 (1-2), 181-238.

Bedendo, Mascia, and Stewart D. Hodges, 2009, The dynamics of the volatility skew: a Kalman filter approach, Journal of Banking & Finance 33 (6), 1156–1165.

Birru, Justin, and Stephen 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, 151–180.

Black, Fischer, and Myron Scholes, 1973, The Pricing of Options and Corporate Liabilities, Journal of Political Economy 81 (3), 637-654.

Bliss, Robert R., and Nikolaos Panigirtzoglou, 2002, Testing the stability of implied probability density functions, Journal of Banking and Finance 26, 381–422.

Bollen, Nicolas P. B., and Robert E. Whaley, 2004, Does Net Buying Pressure Affect the Shape of Implied Volatility Functions?, The Journal of Finance 59 (2), 711-754.

Breeden, Douglas T., and Robert H. Litzenberger, 1978, Prices of State-Contingent Claims Implicit in Option Prices, The Journal of Business 51 (4), 621-651.

Bronstein, I. N., K. A. Semendjajew, G. Musiol, and H. Mühlig (Eds.), 2013, Taschenbuch der Mathematik (Europa-Lehrmittel, Haan-Gruiten).

Chen, Si, Zhen Zhou, and Shenghong Li, 2016, An efficient estimate and forecast of the implied volatility surface: A nonlinear Kalman filter approach, Economic Modelling 58 (C), 655–664.

Chicago Board Options Exchange (CBOE), 2010, The CBOE SKEW INDEX® - SKEW®, URL: https://www.cboe.com/micro/skew/documents/skewwhitepaperjan2011.pdf.

Cont, Rama, and José Da Fonseca, 2002, Dynamics of implied volatility surfaces, Quantitative Finance 2, 45–60.

Corrado, Charles J., and Tie Su, 1996, S&P 500 index option tests of Jarrow and Rudd's approximate option valuation formula, Journal of Futures Markets 16 (6), 611–629.

Corrado, Charles J., and Tie Su, 1997, Implied volatility skews and stock index skewness and kurtosis implied by S&P 500 index option prices, Journal of Derivatives 4, 8–19.

Cox, John C., and Stephen A. Ross, 1976, The Valuation of Options for Alternative Stochastic Processes, Journal of Financial Economics 3, 145-166.

Del Brio, Esther B., Trino Ñiguez, and Javier Perote, 2009, Gram-Charlier densities: a multivariate approach, Quantitative Finance 9 (7), 855–868.

Farkas, Walter, Ciprian Necula, and Boris Waelchli, 2015, Herding and stochastic volatility, Swiss Finance Institute Research Paper No. 15-59.

Figlewski, Stephen, 2010, Estimating the implied risk-neutral density for the U.S. market portfolio, in Tim Bollerslev, Jeffrey R. Russell and Mark Watson, eds.: Volatility and Time Series Econometrics: Essays in Honour of Robert F. Engle (Oxford University Press, Oxford and New York).

Fama, Eugene F., 1965, The Behaviour of Stock-Market Prices, Journal of Business 38 (1), 34-105.

Gan, Guojun, Chaoqun Ma, and Hong Xie, 2014, Measure, Probability and Mathematical Finance: A Problem-Oriented Approach (John Wiley & Sons, Hoboken, New Jersey).

Gemmill, Gordon and Apostolos Saflekos, 2000, How Useful are Implied Distributions? Evidence from Stock Index Options, The Journal of Derivatives 7 (3), 83-91.

Geske, Robert, 1979, The valuation of compound options, Journal of Financial Economics 7, 63-81.

Goncalves, Silvia, and Massimo Guidolin, 2006, Predictable dynamics in the S&P 500 index options implied volatility surface, Journal of Business 79 (3), 1591–1635.

Hafner, Reinhold, and Bernd Schmid, 2005, A Factor-Based Stochastic Implied Volatility Model, Working Paper, risklab germany.

Han, Bing, 2008, Investor Sentiment and Option Prices, Review of Financial Studies 21 (1), 387-414.

Harris, Richard D. F., and C. Coskun Küçüközmen, 2001, The Empirical Distribution of UK and US Stock Returns, Journal of Business Finance & Accounting 28 (5-6), 715-740.

Heston, Steven L., 1993, A Closed Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options, Review of Financial Studies 6 (2), 327-343.

Homescu, Christian, 2011, Implied Volatility Surface: Construction Methodologies and Characteristics, working paper.

Hull, John, and Alan White, 1987, The Pricing of Options on Assets with Stochastic Volatilities, Journal of Finance 42 (2), 281-300.

Jarrow, Robert, and Andrew Rudd, 1982, Approximate option valuation for arbitrary stochastic processes, Journal of Financial Economics 10 (3), 347–369.

Jackwerth, Jens Carsten, 2004, Option-Implied Risk-Neutral Distributions and Risk Aversion (Research Foundation of AIMR, Charlotteville).

Jondeau, Eric, and Michael Rockinger, 2001, Gram-Charlier Densities, Journal of Economic Dynamics and Control 25 (10), 1457-1483.

Jondeau, Eric, and Michael Rockinger, 2000, Reading the smile: the message conveyed by methods which infer risk neutral densities, Journal of International Money and Finance 19, 885–915.

Jurczenko, Emmanuel, Bertrand Maillet, and Bogdan Negrea, 2004, A note on skewness and kurtosis adjusted option pricing models under the Martingale restriction, Quantitative Finance 4 (5), 479-488.

Lebedev, N. N., and Richard A. Silverman (Ed.), 1965, Special Functions and Their Applications (Prentice-Hall, Englewood Cliffs, N.J.).

Ledoit, Oliver, and Pedro Santa-Clara, 1998, Relative Pricing of Options with Stochastic Volatility, University of California-Los Angeles, Finance Working Paper 9-98.

Lynch, Damien, and Nikolaos Panigirtzoglou, 2008, Summary statistics of option-implied probability density functions and their properties, Working Paper No. 345, Bank of England.

Macbeth, James D., and Larry J. Merville, 1979, An Empirical Examination of the Black Scholes Call Option Pricing Model, Journal of Finance 34 (5), 1173-1186.

Malz, Allan M., 2014, A Simple and Reliable Way to Compute Option-Based Risk-Neutral Distributions, Staff Report No. 677, Federal Reserve Bank of New York.

Melick, William R., and Charles P. Thomas, 1997, Recovering an Asset's Implied PDF from Option Prices: An Application to Crude Oil during the Gulf Crisis, Journal of Financial and Quantitative Analysis 32 (1), 91-115.

Merton, Robert C., 1973, Theory of Rational Option Pricing, The Bell Journal of Economics and Management Science 4 (1), 141-183.

Merton, Robert C., 1976, Option Pricing when Underlying Stock Returns are Discontinuous, Journal of Financial Economics 3, 125-144.

Myller-Lebedeff, Wera, 1907, Die Theorie der Integralgleichungen in Anwendung auf einige Reihenentwicklungen, Mathematische Annalen 64, 388-416.

Necula, Ciprian, Gabriel Drimus, and Walter Farkas, 2016, A General Closed Form Option Pricing Formula, Swiss Finance Institute Research Paper, No. 15-53.

Pimbly, Joe, 2006, Gauss-Hermite Quadrature in Financial Risk Analysis, URL: http://www.maxwell-consulting.com/Gauss_Hermite_GARP_Risk_Review_website.pdf.

Rubinstein, Mark, 1983, Displaced Diffusion Option Pricing, Journal of Finance 38 (1), 213-217.

Samuelson, Paul A., 1965, Rational Theory of Warrant Pricing, Industrial Management Review 6 (2), 13-39.

Samuelson, Paul A., and Robert C. Merton, 1969, A Complete Model of Warrant Pricing that Maximizes Utility, Industrial Management Review 10, 17–46.

Schönbucher, Philipp J., 1999, A market model for stochastic implied volatility, Philosophical Transactions of the Royal Society 357 (1758), 2071–2092.

Shiratsuka, Shigenori, 2001, Information Content of Implied Probability Distributions: Empirical Studies of Japanese Stock Price Index Options, Institute for Monetary and Economic Studies, Bank of Japan, November 2001, 142-170.

Skiadopoulos, George, Stewart Hodges, and Les Clewlow, 2000, The dynamics of the S&P 500 implied volatility surface, Review of Derivatives Research 3 (3), 263–282.

8. Appendices

Appendix A – Gauss-Hermite polynomials

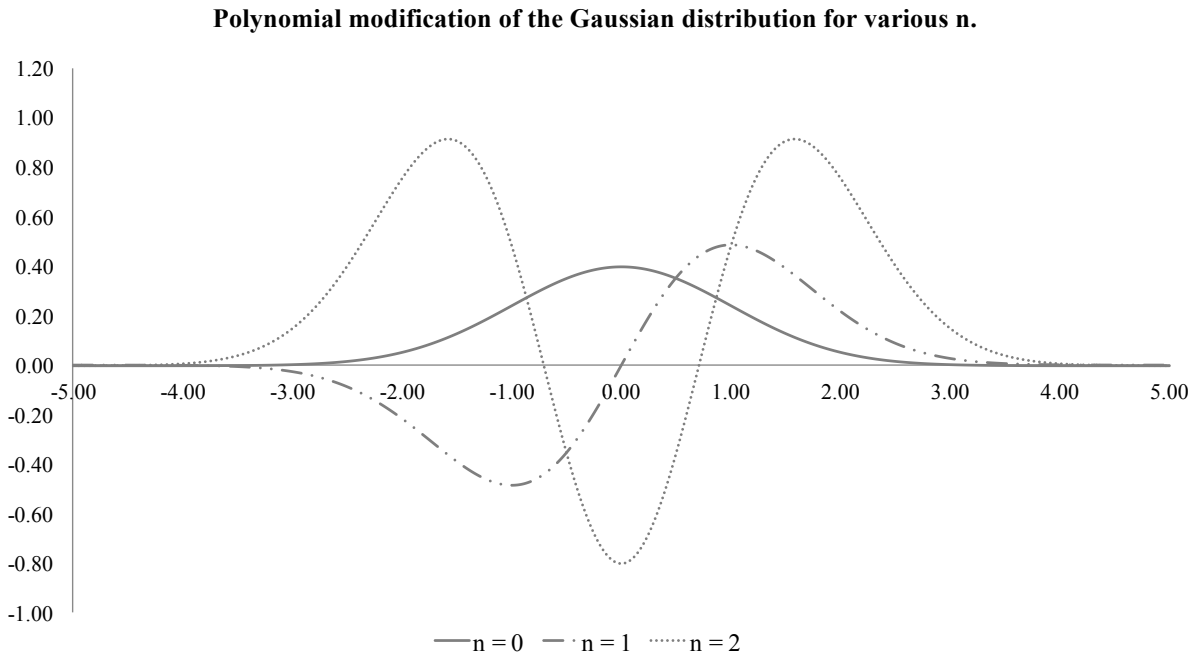

We plot polynomial modifications $z(x)H_n(x)$ of the standard normal distribution $z(x)$ for various orders n, the sum (weighted by GH-expansion coefficients) of which approximates the RND for $n \longrightarrow \infty$.

Appendix B - Minimization problem

Let $K_{t,\tau}$ denote a column vector with length m, containing the exercise prices of all m options traded at time t and with maturity $t+\tau$ (filtered as outlined in *Part 4.1*). Once we obtained estimates for mean and volatility (as described in *Part 3.3*), we can for each element of $K_{t,\tau}$ calculate the corresponding d_1 and d_2 (see Equation (5) in *Part 3.3*). Given this, we can similarly calculate all corresponding I_n and J_n. We then store those values in matrices $I_{t,\tau}$ and $J_{t,\tau}$ of dimension $(m \times 21)$ such that

$$I_{t,\tau} = \begin{bmatrix} I_0(K_1) & \cdots & I_{20}(K_1) \\ \vdots & \ddots & \vdots \\ I_0(K_m) & \cdots & I_{20}(K_m) \end{bmatrix}, J_{t,\tau} = \begin{bmatrix} J_0(K_1) & \cdots & J_{20}(K_1) \\ \vdots & \ddots & \vdots \\ J_0(K_m) & \cdots & J_{20}(K_m) \end{bmatrix}.$$

Now, modifying

$$\tilde{I}_{t,\tau} = S_t e^{\left(\pi_\tau - r + \frac{\sigma_\tau^2}{2}\right)\tau} I_{t,\tau}$$

and

$$\tilde{J}_{t,\tau} = e^{-r\tau} K_{t,\tau} \otimes J_{t,\tau}$$

where $\otimes$ denotes row-wise multiplication. We now can combine

$$C_{t,\tau} = \tilde{I}_{t,\tau} - \tilde{J}_{t,\tau}$$

which gives the estimated prices of each option depending on its strike price by

$$\hat{E}_{t,\tau} = C_{t,\tau} * A_{t,\tau}$$

where $A_{t,\tau}$ is the vector of GH expansion coefficients and $\hat{E}_{t,\tau}$ is the vector of estimated calls. Thus, to find the GH expansion coefficients, we can simply minimize the absolute difference between estimated ($\hat{E}_{t,\tau}$) and observed ($E_{t,\tau}$) prices by finding values for $A_{t,\tau}$ that minimize

$$min \left\| E_{t,\tau} - \hat{E}_{t,\tau} \right\|^2 = min \left\| E_{t,\tau} - C_{t,\tau} * A_{t,\tau} \right\|^2.$$

This yields estimates of $\hat{A}_{t,\tau}$ which can be found in a least-squares sense. However, in the estimation process, one has to consider several constraints to ensure that $\hat{A}_{t,\tau}$ yield a RND that adheres to the properties of a probability measure. These constraints are outlined in *Appendix C*.

Appendix C - Linear constraints

First, to ensure that the probability mass equals one, we must impose

$$\int_{-\infty}^{\infty} p(x)\,dx = \int_{-\infty}^{\infty} \frac{1}{\sigma} z\left(\frac{x-\mu}{\sigma}\right) \sum_{n=0}^{\infty} a_n H_n\left(\frac{x-\mu}{\sigma}\right) dx = 1$$

$$\int_{-\infty}^{\infty} p(x)\,dx = \sum_{n=0}^{\infty} a_n \int_{-\infty}^{\infty} \frac{1}{\sigma} z\left(\frac{x-\mu}{\sigma}\right) H_n\left(\frac{x-\mu}{\sigma}\right) dx = 1.$$

It turns out that this can be simplified to

$$\int_{-\infty}^{\infty} p(x)\,dx = \sum_{n=0}^{\infty} a_n \int_{-\infty}^{\infty} z(x) H_n(x)\,dx = 1.$$

such that we can get the linear constraint

$$\widetilde{H}_{t,\tau} * A_{t,\tau} = 1 \qquad \text{(Constr. 1)}$$

where

$$\widetilde{H}_{t,\tau}{}^{T} = \begin{pmatrix} \int_{-\infty}^{\infty} z(x) H_0(x)\,dx \\ \cdots \\ \cdots \\ \cdots \\ \int_{-\infty}^{\infty} z(x) H_{20}(x)\,dx \end{pmatrix}.$$

It is not straightforward to derive the integrals in $\widetilde{H}_{t,\tau}{}^{T}$ analytically, but they can easily be obtained numerically. This derivation gives some intuition about the nature of the constraint of unitary mass. It is also equivalent to imposing

$$\sum_{k=0}^{\infty} a_{2k} \frac{(2k)!}{k!} = 1$$

truncated such that $2k \leq 20$, which is given in Necula, Drimus and Farkas (2016). Second, to enforce non-negativity on the RND, one must have

$$p(x) = \frac{1}{\sigma} z\left(\frac{x-\mu}{\sigma}\right) \sum_{n=0}^{\infty} a_n H_n\left(\frac{x-\mu}{\sigma}\right) \geq 0$$

for all $x \in (-\infty; \infty)$ and hence

$$\sum_{n=0}^{\infty} a_n z\left(\frac{x-\mu}{\sigma}\right) H_n\left(\frac{x-\mu}{\sigma}\right) \geq 0.$$

This could be derived using the characteristic function, or one can simply introduce small discrete steps on the x-axis and ensure that $p(x) > 0$ at all steps, i.e. that

$$\tilde{X}_{t,\tau} * A_{t,\tau} \geq Z$$

or equivalent

$$-\tilde{X}_{t,\tau} * A_{t,\tau} \leq Z \qquad \text{(Constr. 2)}$$

for

$$\tilde{X}_{t,\tau} = \begin{bmatrix} z\left(\frac{x_i-\mu}{\sigma}\right) H_0\left(\frac{x_i-\mu}{\sigma}\right) & \cdots & z\left(\frac{x_i-\mu}{\sigma}\right) H_{20}\left(\frac{x_i-\mu}{\sigma}\right) \\ \vdots & \ddots & \vdots \\ z\left(\frac{x_n-\mu}{\sigma}\right) H_0\left(\frac{x_n-\mu}{\sigma}\right) & \cdots & z\left(\frac{x_n-\mu}{\sigma}\right) H_{20}\left(\frac{x_n-\mu}{\sigma}\right) \end{bmatrix}$$

and where Z is a vector of zeros. We find that having many linear constraints with x ranging from $x_i = -10$ to $x_n = 10$ and discrete steps of size 0.001 is sufficient to ensure non-negativity in practice. Lastly, the as outlined in Necula, Drimus and Farkas (2016), the martingale restriction is given by

$$e^{\left(\pi_{t,\tau}-(r-q)+\frac{\sigma_{t,\tau}^2}{2}\right)\tau} \sum_{n=0}^{20} a_{n,t,\tau} i^n H_n\left(-i\sigma_{t,\tau}\sqrt{\tau}\right) = 1\,, i = \sqrt{-1} \qquad \text{(Constr. 3)}$$

Thus, all constraints on the RND can be expressed in a linear fashion. Hence, we can use the a linear least-squares optimization (Lsqlin in Matlab) to find $\hat{A}_{t,\tau}$ by

$$\min \left\| E_{t,\tau} - C_{t,\tau} * A_{t,\tau} \right\|^2$$

under *Restrictions (1)*, *(2)* and *(3)*. As outlined in *Part 3.3*, we in practice however only enforce non-negativity and the martingale restriction.

Appendix D - Expressions for skewness and kurtosis

First, for a continuous probability density $f(x)$ and some function $g(x)$, we have

$$E(g(x)) = \int_{-\infty}^{+\infty} g(x)f(x)dx.$$

Now let us define the central moment of m[th] order of the probability density as

$$E(X^m) = E((X - \mu_x)^m).$$

Skewness and kurtosis however are often referred to as normalized central 3rd and 4th moments, i.e.

$$\gamma_1(x) = E(\frac{(X-\mu_x)^3}{\sigma^3}) \; = \; \int_{-\infty}^{+\infty} \frac{(X-\mu_x)^3}{\sigma^3} f(x)dx$$

and

$$\gamma_2(x) = E(\frac{(X-\mu_x)^4}{\sigma^4}) \; = \; \int_{-\infty}^{+\infty} \frac{(X-\mu_x)^4}{\sigma^4} f(x)dx.$$

Now, in our case the probability density function is

$$p(x) = \frac{1}{\sigma} z\left(\frac{x-\mu}{\sigma}\right) \sum_{n=0}^{\infty} a_n \, H_n\left(\frac{x-\mu}{\sigma}\right)$$

and standardizing $\tilde{x} = \frac{x-\mu}{\sigma}$, we have

$$p(\tilde{x}) = \frac{1}{\sigma} z(\tilde{x}) \sum_{n=0}^{\infty} a_n \, H_n(\tilde{x}).$$

Now we could use the integral definition of skewness and kurtosis as above, but the derivation is easier using the characteristics function of the distribution. As defined in Necula, Drimus and Farkas (2016), the characteristic function for the normalized distribution is

$$\varphi(\phi) = e^{-\frac{\phi^2}{2}} \sum_{n=0}^{\infty} a_n \, i^n H_n(\phi), i = \sqrt{-1}.$$

The m[th] moment of a distribution is defined as the m[th] derivative of φ, evaluated at zero, i.e.

$$E(X^m) = i^{-m} \left[\frac{d^m}{d\phi^m} \varphi_X(\phi)\right]_{\phi=0}.$$

Thus,

$$E(X^m) = i^{-m}\left[\frac{d^m}{d\phi^m}e^{-\frac{\phi^2}{2}}\sum_{n=0}^{\infty}a_n\, i^n H_n(\phi)\right]_{\phi=0}.$$

First, we can note that due to the 'summation' rule in calculus, we can evaluate each term of the sum individually. Also, as a_n is a constant, we can write

$$E(X^m) = \sum_{n=0}^{\infty}a_n\, i^{n-m}\left[\frac{d^m}{d\phi^m}e^{-\frac{\phi^2}{2}}H_n(\phi)\right]_{\phi=0}.$$

A useful insight is again, that the expression is linear in a_n, hence we can write

$$E(X^m) = \sum_{n=0}^{\infty}a_n\, D_n^m,\;\; D_n^m = i^{n-m}\left[\frac{d^m}{d\phi^m}e^{-\frac{\phi^2}{2}}H_n(\phi)\right]_{\phi=0},$$

where the expressions for D_n^m will be derived in the following. Now, for the skewness $\gamma_1(x)$, we have $m = 3$, and hence

$$D_n^3 = i^{n-3}\left[\frac{d^3}{d\phi^3}e^{-\frac{\phi^2}{2}}H_n(\phi)\right]_{\phi=0}.$$

For simplicity, let's denote $e^{-\frac{\phi^2}{2}} = E$ and $H_n(\phi) = H_n$ and $\frac{d\,[\,]}{d\phi} = [\,]'$. Thus,

$$D_n^3 = i^{n-3}[EH_n]'''_{\phi=0}$$

$$D_n^3 = i^{n-3}[E'H_n + EH'_n]''_{\phi=0}$$

$$D_n^3 = i^{n-3}[E''H_n + E'H'_n + E'H'_n + EH''_n]'_{\phi=0}$$

$$D_n^3 = i^{n-3}[E'''H_n + E''H'_n + E''H'_n + E'H''_n + E''H'_n + E'H''_n + E'H''_n + EH'''_n]_{\phi=0}$$

Notice, that all terms of E''' and E' equal zero at $\phi = 0$, thus we can drop them completely and have

$$D_n^3 = i^{n-3}[E''H'_n + E''H'_n + E''H'_n + EH'''_n]_{\phi=0} = i^{-3}[3E''H'_n + EH'''_n]_{\phi=0}.$$

Next, the only non-zero terms in E and E'' at $\phi = 0$ are 1 and -1, respectively. Hence,

$$D_n^3 = i^{n-3}[-3H'_n + H'''_n]_{\phi=0}.$$

Now, as H_n is a polynomial, we can evaluate each of its terms individually in the derivative. Since we are considering the third (H'''_n) and first (H'_n) derivative, only terms of H_n with exponents of 3 and 1 will remain, while the others will equal 0 at $\phi = 0$. Since by definition

H_n where n is even do not have those terms, we already know that $D^3_{n=even} = 0$. This is a very intuitive result, since H_n of even order are symmetrical around 0 and do not introduce skewness. It remains to show how D^3_n is defined for odd n. Now, since the non-zero terms in ${H'}_n$ and ${H'''}_n$ will have to be constants (at $\phi = 0$), they thus will be closely related to the coefficients assigned to the x^1 and x^3 terms in H_n. In fact, we can write H_n as a combination of those terms and those that drop out at $\phi = 0$ as

$$H_n = \alpha_n x^3 + \beta_n x^1 + \ldots$$

where α_n and β_n are some coefficients, which will yield

$$D^3_n = i^{n-3}[-3{H'}_n + {H'''}_n]_{\phi=0} = 6\alpha_n - 3\beta_n.$$

Thus, we need to find a formula for coefficients α_n and β_n. In fact, due to the recursive nature of the Hermite polynomials, these will only depend on the order of the polynomial. Following the alternative definition of physicist Hermite polynomials in Lebedev (1965) that reads

$$H_n(x) = \sum_{k=0}^{\langle n/2 \rangle} \frac{(-1)^k n!}{k!\,(n-2k)!}(2x)^{n-2k},$$

where $\langle x \rangle$ denotes the floor function (i.e. the greatest integer less than or equal x), it follows from some simplifications that the coefficients can be expressed as

$$\beta_n = (-1)^{\frac{n-1}{2}} \frac{2n!}{\left(\frac{n-1}{2}\right)!}$$

and

$$\alpha_n = -\frac{(n-1)}{3}\beta_n$$

for odd n. Now we can simply write

$$D^3_n = i^{n-3}[6\alpha_n - 3\beta_n] = i^{n-3}\left[6(-1)\frac{(n-1)}{3}\beta_n - 3\beta_n\right]$$

$$D^3_n = -i^{n-3}\left[6\frac{(n-1)}{3}\beta_n + 3\beta_n)\right] = -i^{n-3}[(2n+1)\beta_n]$$

$$D^3_n = -i^{n-3}(2n+1)(-1)^{\frac{n-1}{2}} \frac{2n!}{\left(\frac{n-1}{2}\right)!}$$

Recognizing, that $-i^{n-3}(-1)^{\frac{n-1}{2}} = -(-1)^{\frac{n-3}{2}}(-1)^{\frac{n-1}{2}} = -(-1)^{n-2} = 1$ for odd n, we finally have

$$\gamma_1(x) = \sum_{n=0}^{\infty} a_n D_n^3$$

where

$$D_n^3 = \begin{cases} (2n+1)\dfrac{2n!}{\left(\frac{n-1}{2}\right)!}, n = odd \\ 0, else \end{cases}.$$

For the kurtosis, we can again start from the definition involving the characteristic function:

$$\gamma_2(x) = E(X^4) = \sum_{n=0}^{\infty} a_n D_n^4, \; D_n^4 = i^{n-4}\left[\frac{d^4}{d\phi^4} e^{-\frac{\phi^2}{2}} H_n(\phi)\right]_{\phi=0}$$

Notice that

$$D_n^4 = i^{n-4}[EH_n]''''_{\phi=0}$$

$$D_n^4 = i^{n-4}[E'''H_n + E''H'_n + E''H'_n + E'H''_n + E''H'_n + E'H''_n + E'H''_n + EH'''_n]'_{\phi=0}$$

$$\begin{aligned} D_n^4 = i^{n-4}[&E''''H_n + E'''H_n{}' + E'''H'_n + E''H''_n + E'''H'_n + E''H''_n + E''H''_n + E'H'''_n \\ &+ E'''H'_n + E''H''_n + E''H''_n + E'H'''_n + E''H''_n + E'H'''_n + E'H'''_n \\ &+ EH''''_n]_{\phi=0}. \end{aligned}$$

Again, $E' = E''' = 0$ at $\phi = 0$. Thus,

$$\begin{aligned} D_n^4 = i^{n-4}[&E''''H_n + E''H''_n + E''H''_n + E''H''_n + E''H''_n + E''H''_n + E''H''_n \\ &+ EH''''_n]_{\phi=0} \end{aligned}$$

$$D_n^4 = i^{n-4}[E''''H_n + 6E''H''_n + EH''''_n]_{\phi=0}$$

Also, the non-zero terms at $\phi = 0$ are $E'''' = 3$, $E'' = -1$ and $E = 1$, which yields

$$D_n^4 = i^{n-4}[3H_n - 6H''_n + H''''_n]_{\phi=0}$$

Recognizing that all terms of the physicist Hermite polynomials of odd order (also for their second and fourth derivative) collapse to zero at $\phi = 0$, we already know that $D_n^4 = 0$ for odd n. Thus, for polynomials of even order we can write

$$H_n = \psi_n x^4 + \omega_n x^2 + \vartheta_n + \cdots$$

and then

$$D_n^4 = i^{n-4}[3\vartheta_n - 12\omega_n + 24\psi_n]_{\phi=0}$$

Similar to the skewness, the kurtosis depends on the coefficients associated with individual terms of the Hermite polynomial of order n. Here again, we can express those as functions of n, i.e.

$$\vartheta_n = \begin{cases} \dfrac{n!}{\left(\frac{n}{2}\right)!}(-1)^{\frac{n}{2}}, n > 1, n = even \\ 0, n = odd \end{cases}$$

$$\omega_n = -n\vartheta_n$$

$$\psi_n = \frac{n(n-2)}{6}\vartheta_n.$$

Combining and recognizing that $-i^{n-4}(-1)^{\frac{n}{2}} = 1$ for all even n, we get

$$D_n^4 = i^{n-4}[3\vartheta_n - 12\omega_n + 24\psi_n]_{\phi=0} = i^{n-4}[3\vartheta_n + 12n\vartheta_n + 4n(n-2)\vartheta_n]_{\phi=0}$$

$$D_n^4 = i^{n-4}[(4n^2 + 4n + 3)\vartheta_n]_{\phi=0}$$

$$D_n^4 = \begin{cases} (4n^2 + 4n + 3)\dfrac{n!}{\left(\frac{n}{2}\right)!}, n = even \\ 0, else \end{cases}$$

and

$$\gamma_2(x) = E(X^4) = \sum_{n=0}^{\infty} a_n D_n^4.$$

An easy way of verifying these expressions for skewness and kurtosis is to use numerical methods (e.g. the integral function in Matlab) to compute values according to the integral definition of those moments. We found that both methods yield the same results. Also, using only $n = 0$, we see that the expressions yield $\gamma_1 = 0$ and $\gamma_2 = 3$, which correspond to the Gaussian.

Appendix E – Tail behaviour of GH-RNDs

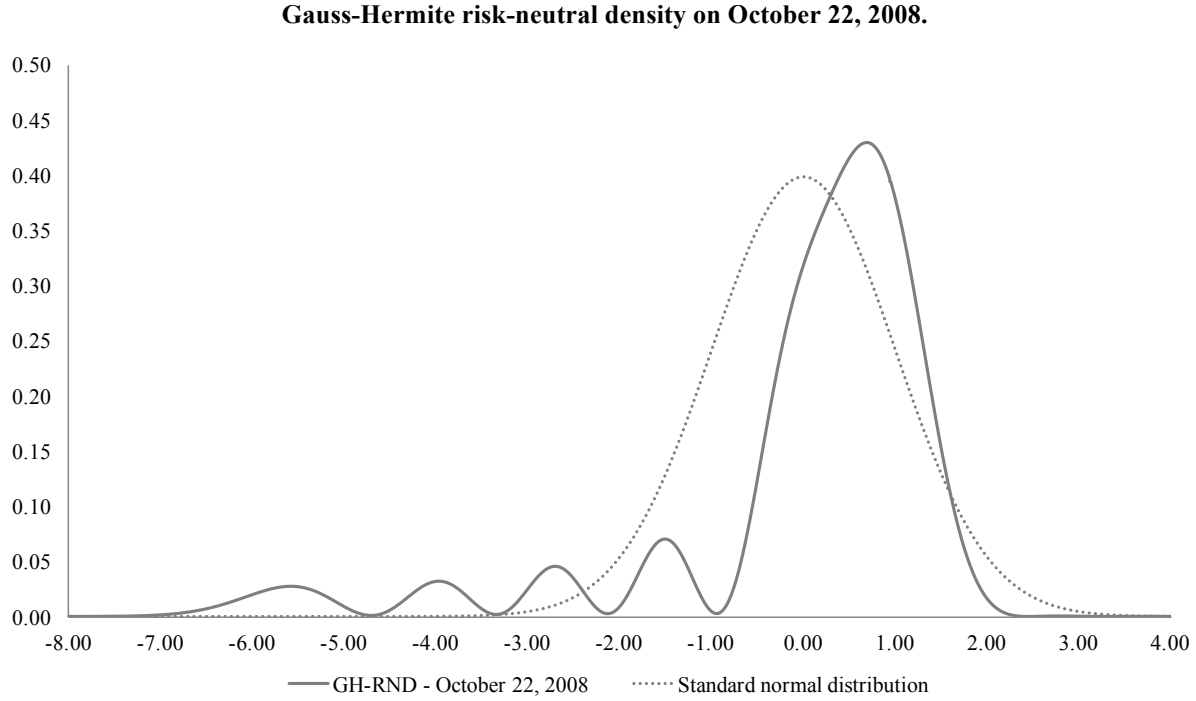

This graph shows the Gauss-Hermite risk-neutral density (standardized) estimated from option quotes on October 22, 2008 and with a time to maturity of $\tau = 0.91$. It also shows a standard normal distribution.

Appendix F – Connection between skewness and kurtosis indices

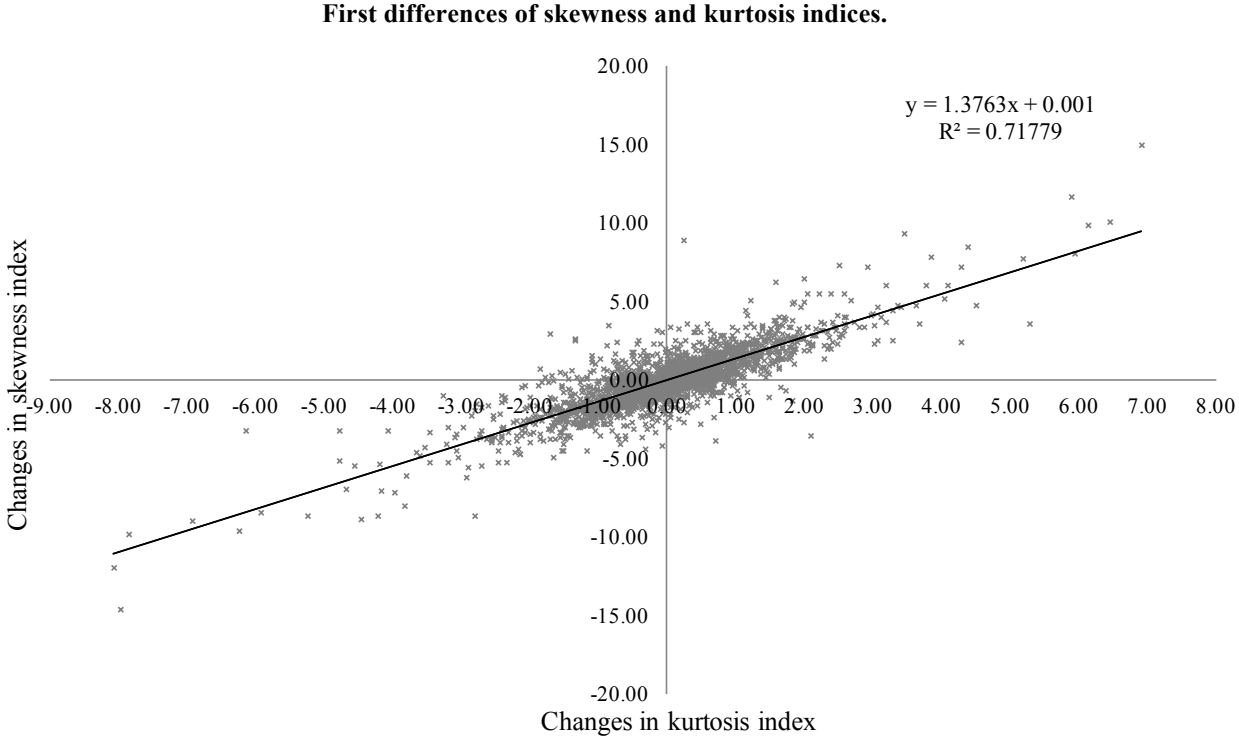

This graph shows the (simultaneous) first differences of the skewness and kurtosis indices.

Appendix G - VAR model output and Granger causality test

Information criteria for VAR models including first differences in the S&P 500 and skewness index (multiple lag-lengths).

lag	AIC	HQIC	SBIC
1	-1.6484	-1.6448	-1.6385
2	-1.6650	-1.6578	-1.6452*
3	-1.6702	-1.6594	-1.6406
4	-1.6699	-1.6555	-1.6303
5	-1.6789	-1.6609*	-1.6295
6	-1.6779	-1.6563	-1.6186
7	-1.6820	-1.6568	-1.6128
8	-1.6844	-1.6556	-1.6053
9	-1.6823	-1.6499	-1.5933
10	-1.6876	-1.6515	-1.5887
11	-1.6859	-1.6463	-1.5771
12	-1.6899	-1.6466	-1.5712
13	-1.6870	-1.6402	-1.5585
14	-1.6859	-1.6355	-1.5475
15	-1.6863	-1.6322	-1.5379
16	-1.6915	-1.6338	-1.5333
17	-1.6927	-1.6314	-1.5246
18	-1.6979*	-1.6331	-1.5200
19	-1.6954	-1.6270	-1.5075
20	-1.6975	-1.6255	-1.4998

This table contains statistical information criteria for a VAR model including both changes of skewness and log-changes in the S&P 500 for up to 20 lags. (*) denotes optimal lag-length according to a given information criterion.

From the above table containing statistical information criteria for the optimal lag-length of vector-autoregressive (VAR) models, we see that the optimal number of terms to be included in the model is not clear, as each criterion suggest a different lag-length. We suggest using 5 lags, as it balances the other criteria and is easy to interpret (5 days usually is one week of trading). The estimated VAR model then will take the following form:

$$\begin{bmatrix} \Delta S_t \\ \Delta\bar{\gamma}_{1,t} \end{bmatrix} = \begin{bmatrix} a_{1,1}^1 & a_{1,2}^1 \\ a_{2,1}^1 & a_{2,2}^1 \end{bmatrix} \begin{bmatrix} \Delta S_{t-1} \\ \Delta\bar{\gamma}_{1,t-1} \end{bmatrix} + \ldots + \begin{bmatrix} a_{1,1}^5 & a_{1,2}^5 \\ a_{2,1}^5 & a_{2,2}^5 \end{bmatrix} \begin{bmatrix} \Delta S_{t-5} \\ \Delta\bar{\gamma}_{1,t-5} \end{bmatrix} + \begin{bmatrix} e_{1,t} \\ e_{2,t} \end{bmatrix}$$

where ΔS_t and $\Delta\bar{\gamma}_{1,t}$ denote the S&P 500 log-return and the first difference of the skewness index, respectively. The model output is shown below.

VAR model output.

Equation	Parameters	RMSE	R-squared	Chi-squared	p-value
ΔS_t	10	0.0134	0.0252	60.4589	0
$\Delta\bar{\gamma}_{1,t}$	10	1.8744	0.0861	220.7700	0

Equation:	ΔS_t				
Variable	Lag	Coefficient	Std. Err.	z	P>(z)
ΔS_{t-1}	1	-0.1144*	0.0207	-5.5300	0.0000
ΔS_{t-2}	2	-0.0630*	0.0208	-3.0300	0.0020
ΔS_{t-3}	3	0.0246	0.0208	1.1800	0.2370
ΔS_{t-4}	4	-0.0253	0.0208	-1.2200	0.2240
ΔS_{t-5}	5	-0.0519	0.0206	-2.5100	0.0120
$\Delta\bar{\gamma}_{1,t-1}$	1	0.0003	0.0001	1.9100	0.0560
$\Delta\bar{\gamma}_{1,t-2}$	2	0.0003	0.0002	1.8600	0.0620
$\Delta\bar{\gamma}_{1,t-3}$	3	-0.0003	0.0002	-1.6900	0.0910
$\Delta\bar{\gamma}_{1,t-4}$	4	-0.0003	0.0002	-1.8800	0.0600
$\Delta\bar{\gamma}_{1,t-5}$	5	-0.0001	0.0001	-0.7000	0.4860

Equation:	$\Delta\bar{\gamma}_{1,t}$				
Variable	Lag	Coefficient	Std. Err.	z	P>(z)
ΔS_{t-1}	1	-6.2651	2.8898	-2.1700	0.0300
ΔS_{t-2}	2	3.9980	2.9079	1.3700	0.1690
ΔS_{t-3}	3	1.1297	2.9137	0.3900	0.6980
ΔS_{t-4}	4	1.5488	2.9047	0.5300	0.5940
ΔS_{t-5}	5	0.9281	2.8851	0.3200	0.7480
$\Delta\bar{\gamma}_{1,t-1}$	1	-0.2721*	0.0206	-13.2000	0.0000
$\Delta\bar{\gamma}_{1,t-2}$	2	-0.1455*	0.0214	-6.8100	0.0000
$\Delta\bar{\gamma}_{1,t-3}$	3	-0.1123*	0.0215	-5.2300	0.0000
$\Delta\bar{\gamma}_{1,t-4}$	4	-0.0663*	0.0214	-3.1000	0.0020
$\Delta\bar{\gamma}_{1,t-5}$	5	-0.0975*	0.0206	-4.7300	0.0000

This table shows the estimated VAR model output for both changes in the skewness index ($\Delta\bar{\gamma}_{1,t}$) and log-changes in the underlying (ΔS_t). Statistically significant coefficients (at 1%) are denoted by (*).

Granger causality Wald test.

Equation	Excluded	Chi-squared	P>(Chi-squared)
ΔS_t	$\Delta\bar{\gamma}_{1,t}$	14.482	0.013
$\Delta\bar{\gamma}_{1,t}$	ΔS_t	7.6969	0.174

This table shows the results of a Wald test for Granger causality. It tests the null-hypothesis that coefficients associated with all lags of the excluded variable are zero in the given equation of the VAR model.