Predictable dynamics in higher order risk-neutral moments: Evidence from the S&P 500 options*

Michael Neumann[a] and George Skiadopoulos[b]

This Draft: 18/05/2011

Abstract

We investigate whether there are predictable patterns in the dynamics of higher order risk-neutral moments extracted from the S&P 500 options. To this end, we conduct a horse race among alternative forecasting models within an out-of-sample context. We consider both a statistical and an economic setting. We detect the set of statistically superior models by using the newly developed model confidence set methodology of Hansen, Lunde, and Nason (2010). We assess the economic significance of the forecasts by means of skewness and kurtosis option trading strategies. We find that the risk-neutral higher moments can be statistically forecasted by means of a set of models even over the recent subprime crisis period. In addition, trading implied skewness yields significant risk-adjusted profits. However, this economic significance vanishes once we incorporate transaction costs. The results are robust across various maturities and have implications for the predictability of implied volatility surfaces, development of option pricing models and efficiency of the S&P 500 options market.

Keywords: Implied skewness, Implied kurtosis, Implied volatility surface, Model confidence set, Option strategies, Risk-neutral moments.

JEL Classification: C53, C58, G10, G13, G17

* We would like to thank Turan Bali, Alejandro Bernales, Tim Bollerslev, Robert Engle, Daniel Giamouridis, Massimo Guidolin, Eirini Konstantinidi, Marc Paolella, Efthymios Tsionas and Greg Vilkov for useful discussions and comments. Any remaining errors are our responsibility alone.

[a] HVB-Institute for Mathematical Finance, Munich University of Technology, michael_neumann_ma@yahoo.de Business Campus Munich, Parkring 11, 85748 Garching-Hochbrueck, Germany.

[b] Corresponding author. Department of Banking and Financial Management, University of Piraeus, Greece, Financial Options Research Centre, Warwick Business School, University of Warwick, and Cass Business School, City University, UK gskiado@unipi.gr Postal Address: University of Piraeus, Department of Banking and Financial Management, Karaoli and Dimitriou 80, Piraeus 18534, Greece, Tel: +30-210-4142363.

1. Introduction

At any given point in time, market option prices convey information about the moments of the risk-neutral probability density function (PDF; also termed implied PDF) of the underlying asset; the horizon of the PDF corresponds to the expiry of the options under consideration.[1] Risk-neutral moments (RNMs hereafter) can be calculated either directly (Bakshi, Kapadia and, Madan, 2003, BKM hereafter) or by extracting first the whole implied PDF as suggested by Breeden and Litzenberger (1978, see also Jackwerth, 2004, for an excellent review of the related literature).[2] Even though there is an extensive literature on the dynamics of the second RNM (i.e. the implied volatility, see e.g., Harvey and Whaley, 1992, Konstantinidi, Skiadopoulos and Tzagkaraki, 2008, Goyal and Saretto, 2009, among others), surprisingly little attention has been paid to the evolution of risk-neutral skewness and kurtosis over time per se. This paper fills this void by undertaking a comprehensive study of the dynamics of the S&P 500 higher RNMs and in particular whether these can be used to forecast their subsequent changes.[3]

Understanding the dynamics of the higher order RNMs is of importance to both academics and practitioners since their evolution reflects the evolution of market option prices and in particular that of implied volatility surfaces (IVSs). At any given point in time, an IVS is defined to be the function that maps implied volatilities to different strike prices and maturities (see e.g., Canina and Figlewski, 1993, Gonçalves and Guidolin, 2006, Bernales and Guidolin, 2010, and references therein for the empirical evidence on IVSs). It is well documented that IVSs evolve over time in a complex fashion; their dynamics are driven by two or three factors that affect their level, slope and curvature (see e.g., Skiadopoulos, Hodges and Clewlow, 1999, Fengler, Härdle, and Villa, 2003, Fengler, Härdle, and Mammen, 2007, Carverhill, Cheuk, and Dyrting, 2009, among others, and Alexander, 2008, for a review). Equivalently, any change in the RNMs maps to a change in the IVS, too. This is because changes in risk-neutral volatility, skewness and kurtosis are related to each one of these three factors, respectively (Zhang and Xiang, 2008). Therefore, the study of whether the evolution of

---

[1] The risk-neutral PDF is defined via the relationship that dictates that the market option price equals its theoretical price calculated as the integral of the option payoff with respect to the risk-neutral PDF.

[2] There is also a number of pseudo measures of the respective risk neutral moments proposed by the literature (see e.g., Lynch and Panigirtzoglou, 2008, Rehman and Vilkov, 2010, Xing, Zhang and Zhao, 2010, and Giamouridis and Skiadopoulos, 2010, for a review). It should be noted though that these alternative measures are no moments estimates in a statistical sense but rather measure the concepts underlying the statistical moments (i.e. uncertainty for volatility, asymmetry for skewness, and heavy tails for kurtosis). Therefore, they are not considered for the purposes of the subsequent analysis.

[3] This question is distinct from the question whether RNMs and market option prices in general are useful to address asset allocation (e.g., DeMiguel, Plyakha, Uppal, and Vilkov, 2011, Kostakis, Panigirtzoglou, and Skiadopoulos, 2011), risk management (e.g., Chang, Christoffersen, Jacobs, and Vainberg, 2009) and stock return forecasting (e.g., Bali and Hovakiminiam, 2009, Ang, Bali, and Cakici, 2010, Bakshi, Panayotov and Skoulakis, 2010, Chang, Christoffersen, and Jacobs, 2010, Cremers and Weinbaum, 2010, Rehman and Vilkov, 2010, Xing, Zhang and Zhao, 2010) questions; see also Giamouridis and Skiadopoulos (2010) for an extensive review.

higher order RNMs is predictable amounts to investigating whether the evolution of IVSs is predictable per se. The latter is important for option pricing and trading purposes. In particular, any option pricing model to be credible has not only to price options accurately across the cross-section of options at a given point in time but also to do so at different points in time (see e.g., Bakshi, Cao and Chen, 1997, Dumas, Fleming and Whaley, 1998). To this end, a number of option pricing models try to capture the evolution of IVS via modelling the factors that drive its evolution (e.g., correlation between the underlying asset and volatility and/or time varying jump intensities) and hence implicitly modelling the RNMs (see e.g., Christoffersen, Heston and Jacobs, 2009, Christoffersen, Jacobs and Ornthanalai, 2010, and references therein). Understanding the dynamics of RNMs will shed light on the specification that this type of models should use in order to price and hedge options accurately over different points in time. Furthermore, from a trader's point of view, forecasting changes in the level, slope and curvature of the IVS, i.e. the changes in RNMs moments, can help developing profitable volatility, skewness and kurtosis trading strategies; the prior literature has mostly studied volatility trading strategies (see e.g., Konstantinidi, Skiadopoulos, and Tzagkaraki, 2008, and references therein); to the best of our knowledge, Bali and Murray (2011) is the only paper that look at risk-neutral skewness and kurtosis strategies by exploiting their dynamics.[4]

The evidence on the predictability of IVSs is mixed depending on the evaluation metric and employed dataset. Dumas, Fleming and Whaley (1998) find that the dynamics of the S&P 500 IVS are highly unstable under an options hedging setting. On the other hand, Gonçalves and Guidolin (2006) and Chalamandaris and Tsekrekos (2010) find that the S&P 500 and currency options IVSs are statistically predictable, respectively, whereas Bernales and Guidolin (2010) find that the predictability of equity options IVS is both statistically and economically significant. All these studies fit parametric models to implied volatility surfaces. Instead, studying whether there are predictable patterns in the evolution of RNMs provides an alternative way of investigating whether changes in the IVSs can be forecasted and hence extends the literature on the predictability of IVSs. Furthermore, in the case where RNMs prove to be predictable this will shed light on the factor(s) that the predictability of volatility surfaces may be attributed to.[5]

---

[4] Ait-Sahalia,Wang, and Yared (2001) and Jha and Kalimipalli (2010) also develop higher order RNMs strategies, yet their strategies are based on forecasts for the *realised* moments rather than forecasts for the implied moments per se, as we do in this paper.

[5] Understanding the process that drives RNMs may also be informative for the process that drives the corresponding physical moments. This is helpful for the further development of the literature that deals with the modelling of higher order moments (see e.g., Hansen, 1994, Harvey and Siddique, 1999, Jondeau and Rockinger, 2003 and 2009, Haas, Mittnik, and Paolella, 2004, Brooks, Burke, Heravi and Persand, 2005, Alexander and Lazar, 2006, Kim, Manganelli, and White, 2008). The latter is a prerequisite for the development of dynamic portfolio choice models that take into account the presence of higher moments by postulating a stochastic process for their evolution over time (see e.g., Patton, 2004, and Jondeau and Rockinger, 2008, for such an approach). There is a caveat though in extending the findings on the time

There is already some literature that has studied the factors that affect higher RNMs. Dennis and Mayhew (2002) explore the drivers of risk-neutral skewness of individual stocks extracted from equity option prices. They find that the returns' risk-neutral distributions of stocks with higher betas tend to be more negatively skewed. In addition, firm specific factors such as firm size and trading volume help explaining variation in risk-neutral skewness. Taylor, Yadav, and Zhang (2009) extend this research approach by introducing additional explanatory variables. Hansis, Schlag, and Vilkov (2010) perform a similar analysis by investigating also the determinants of risk-neutral kurtosis. They find that accounting variables such as market-to-book ratio and return on equity, as well as, stock specific factors and measures for the quality of information available to the investor play an important role in explaining variation in risk-neutral skewness and kurtosis. However, these papers focus on studying the cross-sectional properties of risk-neutral moments; little attention has been paid to analysing their dynamics over time and in particular whether they can be predicted. The aim of the current paper is to fill this void. To the best of our knowledge, the papers by Panigirtzoglou and Skiadopoulos (2004), Lynch and Panigirtzoglou (2008), and Hansis, Schlag, and Vilkov (2010) are the closest to ours. The first postulates a process that drives the dynamics of the S&P 500 implied PDF and identifies the relevant factors.[6] The second provides a detailed description of the evolution of RNMs extracted from index and interest rate options. Among other contributions, the authors examine whether cross-correlations between moments may be useful to forecast the future moment values. Similarly, the third paper employs a vector autoregressive model for the RNMs of individual U.S. stocks and finds that it explains the behaviour of the RNMs satisfactorily. However, in all above mentioned papers, the analysis is performed only within sample. Moreover, no other models are considered for comparison purposes.

Instead, we take a more general approach and make at least four contributions to the existing literature. First, we conduct a horse race among a number of models to investigate whether the evolution of higher order RNMs can be forecasted; the dynamics of risk-neutral volatility are also investigated so that to be contrasted with these of the higher-order moments. This is because the question of predictability is always a joint hypothesis test of the forecasting model that is used. Second, we evaluate the *out-of-sample* forecasting performance of the models. To this end, we use a

---

series properties of RNMs to these of the corresponding physical moments. This is because RNMs are related to the respective moments of the physical distribution of the underlying asset's returns via the risk aversion of the representative agent (see e.g., BKM, 2003, Bliss and Panigirtzoglou, 2004, Jackwerth, 2004, Kostakis, Panigirtzoglou, and Skiadopoulos, 2011). In the case that this relationship may be highly non-linear such an extension of results may not be straightforward. Nevertheless, the study of the time series properties of RNMs can still serve as a first step towards exploring these of the corresponding physical moments.

[6] Figlewski (2010) and Birru and Figlewski (2010) can also be regarded as related papers since they examine the relation between the quantiles of the S&P implied PDF and the S&P 500.

number of performance measures and apply formal statistical tests to evaluate the statistical significance of these measures. Furthermore, we employ the newly developed model confidence set methodology proposed by Hansen, Lunde, and Nason (2010) to determine the best performing models. Third, we assess the economic significance of the forecasts of the RNMs derived by each forecasting model by implementing option trading strategies based on the statistical forecasts for the future values of the RNMs. To this end, we propose and apply skewness and kurtosis option trading strategies extending the evidence on the performance of higher order moment strategies provided by Bali and Murray (2011). Finally, to address the posed question, we employ a dataset of European options written on the S&P 500 spanning the period from 1996 to 2009. This is a rich period that includes bull and bear regimes as well as the recent subprime crisis period; the use of such a rich period is necessary to uncover the dynamics of RNMs (Broadie, Chernov, and Johannes, 2007). To check the robustness of our results, we examine higher RNMs of alternative constant horizons. In addition, we study the effect of the recent subprime crisis.

The remainder of this paper is organised as follows: Section 2 describes the dataset. Section 3 describes briefly the BKM method and the way we implement it to extract the RNMs. The next section describes the forecasting models we use to examine the predictability of RNMs. Section 5 presents the in-sample evidence on the forecasting performance of these models and Section 6 assesses the statistical significance of their out-of-sample performance. Sections 7 and 8 investigate the economic significance of the formed forecasts and the effect of the subprime crisis period, respectively. Section 9 concludes, discusses the implications of the results, and suggests directions for future research.

2. Data

We obtain daily S&P 500 European style option data from the Ivy DB database of OptionMetrics from January 1996 to October 2009. For the purposes of our analysis, we use the S&P 500 implied volatilities provided by Ivy DB for each traded contract (see also Engle and Mistry, 2008, and Chang, Christoffersen, Jacobs, and Vainberg, 2009, for a similar approach). Ivy DB provides implied volatilities extracted from options with different strikes and maturities ranging from five to 270 days.[7] It calculates them based on the midpoint of bid and ask prices using Merton's (1973) model to

[7] These options cover on average about five maturities per day with about 40 strikes trading on average per maturity. The number of traded maturities and strikes grows considerably over the sample period as time goes by. The average number of covered maturities per day is five in January 1996 and increases to 10 in October 2009. The average number of strikes per maturity grows from about 24 in January 1996 to about 64 in October 2009. Notice that Ivy DB also provides implied volatility surfaces measured across various levels of delta for fixed maturities. However, we do not use these and prefer using the provided implied volatilities for the respective contracts so that we are able to control the filtering and interpolation mechanics.

account for the fact that the employed options are European and the underlying asset pays dividends (see also the Ivy DB reference manual for a detailed description of the calculation of implied volatilities). In addition, the closing price of the S&P 500 and the continuously paid dividend yield are also obtained from Ivy DB.

Following the existing literature on calculating option-implied distributions/moments, we impose several filters on the option dataset prior to extracting RNMs. First, we consider only out-of-the money (OTM) and at-the-money (ATM) options. Second, we incorporate only options with non-zero bid prices and premiums, measured as the midpoint of best bid and best offer, greater than 3/8 $. We discard options with implied volatilities greater than 100% as well as options for which Ivy DB does not provide implied volatilities. In addition, we remove options with zero open interest and zero trading volume. We also discard options violating Merton's (1973) arbitrage bounds; the continuously compounded dividend yield provided by Ivy DB is used to proxy the continuously compounded dividend yield. Finally, we exclude options that form vertical and butterfly spreads with negative prices (see Jackwerth and Rubinstein, 1996, for a similar approach). In particular, in the case where the value of any one of these spreads is negative, all options involved in the spread are discarded.

U.S. LIBOR rates for maturities of one to six months are obtained from Bloomberg to proxy the risk-free rate; the discretely compounded quotes are converted to their continuously compounded counterparts. We obtain rates for any other required maturity by interpolating linearly across surrounding maturities. In the cases where the desired maturity is beyond the quoted maturities, we linearly extrapolate using the closest available LIBOR rate. Finally, we use data for a set of additional economic variables. The set of these variables consists of the USD/EURO exchange rate, WTI crude oil price, slope of the yield curve computed as the difference of the 10-year government bond and the LIBOR rate, and the volume of the nearest maturity futures contract written on the S&P 500. The government bond data are obtained from the website of the St. Louis Fed and all other economic variables data from Datastream.

3. Extracting risk-neutral moments

We extract RNMs from option prices using the model-free methodology suggested by BKM. This method has been recently applied by number of authors (see e.g., Dennis and Mayhew, 2002, Chang, Christoffersen, Jacobs, and Vainberg, 2009, Chang, Christoffersen and Jacobs, 2010, Hansis, Schlag, and Vilkov, 2010, among others) since it extracts RNMs from option prices without imposing any specific assumptions on the underlying's stochastic process.

3.1 The BKM method

We provide a short description of the BKM method. Let E_t^Q denote the conditional expected value operator under the risk-neutral measure formed at time *t*, *r* the risk-free rate, $C(t,\tau;K)$ ($P(t,\tau;K)$) the price of a call (put) option with time to expiration τ and strike price K and $R(t,\tau)=ln[S(t+\tau)]-ln[S(t)]$ be the continuously compounded rate of return at time *t* over a time period τ. Let also $V(\cdot)$, $W(\cdot)$ and $X(\cdot)$

$$V(t,\tau)\equiv E_t^Q\left\{e^{-r\tau}R(t,\tau)^2\right\} \tag{1}$$

$$W(t,\tau)\equiv E_t^Q\left\{e^{-r\tau}R(t,\tau)^3\right\} \tag{2}$$

$$X(t,\tau)\equiv E_t^Q\left\{e^{-r\tau}R(t,\tau)^4\right\} \tag{3}$$

denote the fair values of three respective contracts with corresponding payoff functions $H[S]$

$$H[S]=\begin{cases} R(t,\tau)^2 \\ R(t,\tau)^3 \\ R(t,\tau)^4 \end{cases} \tag{4}$$

Let $\mu(t,\tau)\equiv E_t^Q\left\{ln\left[\frac{S(t+\tau)}{S(t)}\right]\right\}\approx e^{r\tau}-1-\frac{e^{r\tau}}{2}V(t,\tau)-\frac{e^{r\tau}}{6}W(t,\tau)-\frac{e^{r\tau}}{24}X(t,\tau)$ be the mean of the log-return over period τ. The risk-neutral volatility (*MFIV*), skewness (*SKEW*) and kurtosis (*KURT*) extracted at time *t* with horizon τ can be expressed in terms of the fair values of the three artificial contracts, i.e.

$$MFIV(t,\tau)=\sqrt{E_t^Q\left\{R(t,\tau)^2\right\}-\mu(t,\tau)^2}=\sqrt{V(t,\tau)e^{r\tau}-\mu(t,\tau)^2} \tag{5}$$

$$\begin{aligned} SKEW(t,\tau)&=\frac{E_t^Q\left\{(R(t,\tau)-E_t^Q\left[R(t,\tau)\right])^3\right\}}{\left\{E_t^Q(R(t,\tau)-E_t^Q\left[R(t,\tau)\right])^2\right\}^{\frac{3}{2}}} \\ &=\frac{e^{r\tau}W(t,\tau)-3\mu(t,\tau)e^{r\tau}V(t,\tau)+2\mu(t,\tau)^3}{\left[e^{r\tau}V(t,\tau)-\mu(t,\tau)^2\right]^{\frac{3}{2}}} \end{aligned} \tag{6}$$

$$KURT(t,\tau)=\frac{E_t^Q\left\{(R(t,\tau)-E_t^Q\left[R(t,\tau)\right])^4\right\}}{\left\{E_t^Q(R(t,\tau)-E_t^Q\left[R(t,\tau)\right])^2\right\}^2}$$
$$=\frac{e^{r\tau}X(t,\tau)-4\mu(t,\tau)e^{r\tau}W(t,\tau)+6e^{r\tau}\mu(t,\tau)^2V(t,\tau)-3\mu(t,\tau)^4}{\left[e^{r\tau}V(t,\tau)-\mu(t,\tau)^2\right]^2} \quad (7)$$

BKM show that the arbitrage-free prices of $V(t,\tau),W(t,\tau)$ and $X(t,\tau)$ are given by

$$V(t,\tau)=\int_{S(t)}^{\infty}\frac{2\left(1-ln\left[\frac{K}{S(t)}\right]\right)}{K^2}C\left(t,\tau;K\right)dK+\int_{0}^{S(t)}\frac{2\left(1+ln\left[\frac{S(t)}{K}\right]\right)}{K^2}P\left(t,\tau;K\right)dK \quad (8)$$

$$W(t,\tau)=\int_{S(t)}^{\infty}\frac{6\,ln\left[\frac{K}{S(t)}\right]-3\left(ln\left[\frac{K}{S(t)}\right]\right)^2}{K^2}C\left(t,\tau;K\right)dK-\int_{0}^{S(t)}\frac{6\,ln\left[\frac{S(t)}{K}\right]+3\left(ln\left[\frac{S(t)}{K}\right]\right)^2}{K^2}P\left(t,\tau;K\right)dK \quad (9)$$

$$X(t,\tau)=\int_{S(t)}^{\infty}\frac{12\left(ln\left[\frac{K}{S(t)}\right]\right)^2-4\left(ln\left[\frac{K}{S(t)}\right]\right)^3}{K^2}C\left(t,\tau;K\right)dK$$
$$+\int_{0}^{S(t)}\frac{12\left(ln\left[\frac{S(t)}{K}\right]\right)^2+4\left(ln\left[\frac{S(t)}{K}\right]\right)^3}{K^2}P\left(t,\tau;K\right)dK \quad (10)$$

Thus, the price of each contract can be computed as a linear combination of OTM call and put options. Based on these prices, the risk-neutral volatility, skewness and kurtosis are computed in a model-free manner.

3.2 Empirical implementation

The implementation of equations (8), (9), and (10) requires a continuum of OTM call and put options across strikes. However, market option quotes are available only for a bounded range of discrete strike prices. This will incur a bias in the calculation of RNMs (see e.g. Dennis and Mayhew, 2002, and Jiang and Tian, 2005, 2007). In addition, for the purposes of our study, *constant maturity* moments need to be extracted so that to eliminate the effect of the shrinking time to maturity on the dynamics of the risk-neutral moments as time goes by.

To address both issues, once we apply the data filters described in Section 2 to any given day, we extract the expirations for which there are at least two OTM puts and two OTM calls. We discard maturities that do not satisfy this requirement from the analysis. Then, we convert the strike prices of the remaining options with a given maturity into call deltas using Merton's (1973) model. Subsequently, for any given traded maturity, we interpolate across the implied volatilities to obtain a continuum of implied volatilities in the delta space. Interpolating in the delta rather than the strike price space has been introduced by Malz (1997) and offers two advantages. First, away-from-the-money options are grouped more closely together than near-the-money options. This allows for greater variation in shape near the center of the distribution where the data are more reliable (see Bliss and Panigirtzoglou, 2004). Second, the entire range of strike prices available for each maturity can be used as delta is bounded between zero and one.[8] To construct the constant maturity moments, we proceed in three steps following the approach suggested by Bliss and Panigirtzoglou (2002) and also applied by Bliss and Panigirtzoglou (2004), Panigirtzoglou and Skiadopoulos (2004), Lynch and Panigirtzoglou (2008), and Kostakis, Panigirtzoglou, and Skiadopoulos (2011). First, we choose nine delta values (0.1, 0.2,...,0.9) and for any given delta, we interpolate across volatilities in the time dimension by using a cubic smoothing spline. Then, we select the implied volatility for a targeted expiration; 30, 60 and 90-days expirations are chosen so that to check the effect of the time horizon to the subsequent analysis. Next, we obtain the constant maturity implied volatility curve by fitting a cubic spline through the nine constant maturity implied volatilities. If the target expiration is below the smallest available expiration traded in the dataset, a constant maturity implied volatility curve is not constructed.[9] Finally, we convert the delta grid of the constant maturity implied volatility curve and the corresponding implied volatilities to the associated strike and option prices, respectively, via Merton's (1973) model.[10] Then, in line with Dennis and Mayhew (2002), we compute the constant

[8] We interpolate on a delta grid with 1,000 grid points ranging from 0.01 to 0.99 using a cubic smoothing spline (with smoothing parameter 0.99). Deltas are calculated by using the at-the-money (ATM) implied volatility. This is done to ensure that the ordering of deltas is the same as the ordering of the strike prices. Using the implied volatilities extracted from the respective options could change the ordering of the deltas in cases where the implied volatility curve is steep. Options with deltas above 0.99 and below 0.01 are discarded as these correspond to far OTM options that are not actively traded. We also make sure that for each maturity there are options with deltas below 0.25 and above 0.75 in order to span a wide range of moneyness regions. If this requirement is not satisfied, we discard the respective maturity from the sample. For deltas beyond the largest and the smallest available delta, we extrapolate horizontally using the respective boundary implied volatility (see also among others, Bliss and Panigirtzoglou, 2004, Panigirtzoglou and Skiadopoulos, 2004, Lynch and Panigirtzoglou, 2008, Ang, Bali, and Cakici, 2010, and Kostakis, Panigirtzoglou, and Skiadopoulos, 2011, for a similar use of the delta space). We have also used the smooth pasting condition proposed by Jiang and Tian (2007) for extrapolation purposes. However, this delivers implausibly large implied volatilities.

[9] Extrapolating in the time dimension domain yields time series of implied moments that exhibit artificially created spikes.

[10] The use of Merton's (1973) model for the translation of implied volatilities into strike prices is consistent with the approach taken by Ivy DB. The model serves only as translation mechanisms from market option prices to implied volatilities; its use does not imply that options are correctly priced by it.

maturity moments [equations (5), (6), (7)] by evaluating the integrals in formulae (8), (9), and (10) using trapezoidal approximation.

3.3 Risk-neutral moments: Results

Figure 1 shows the evolution of the 30, 60, and 90-days constant maturity S&P 500 risk-neutral moments over the sample period. The risk-neutral volatility is annualised based on the day count convention moments maturity/365 which is consistent with the CBOE VIX convention. We can see that the risk-neutral skewness is negative throughout the entire sample period and for any given maturity. In addition, we can see that in general there is excess kurtosis. These findings are consistent with the empirical evidence on left skewed heavy-tailed risk-neutral distributions (e.g., Jackwerth, 2004) that may be due to factors such as stochastically varying volatility and/or occurrence of extreme events (see e.g., Bates, 2003). Notice that in line with the findings of Ait-Sahalia and Lo (1998) and Carr and Wu (2003), the skewness and kurtosis do not converge to zero and three, respectively, as the horizon increases as one would expect by the central limit theorem.

Table 1 reports the descriptive statistics of the extracted S&P 500 moments (mean, median, standard deviation, skewness, and kurtosis) for the three maturities over the sample period from 1996 to 2009. The percentage of observations (NA percentage) excluded due to filters applied to data is also reported. We can see that the means of risk-neutral skewness and kurtosis are negative and greater than three, respectively, for all maturities. Both findings are consistent with prior literature that documents that the index risk-neutral distributions are negatively skewed and exhibit excess kurtosis (see e.g., BKM, 2003, Chang, Christoffersen, Jacobs, and Vainberg, 2009, Figlewski, 2010, among others).

For the purposes of our analysis, we divide the sample into an in-sample part from January 1996 to January 2000, and an out-of-sample part spanning the remainder of the dataset. Table 2 reports the summary statistics of the S&P 500 risk-neutral moments measured in levels (Panel A) and first differences (Panel B) over the in-sample period. The first order autocorrelation is reported for the levels and first differences series. The values of the augmented Dickey Fuller (ADF) test statistic are also reported. One and two asterisks denote rejection of the null hypothesis (zero first-order autocorrelation, existence of a unit root) at 5% and 1% levels, respectively. We can see that every moment is positively autocorrelated in the levels, while the skewness and kurtosis is negatively autocorrelated in the first differences. In addition, the unit root test results reveal that all moments (but the 30 days implied kurtosis) are integrated of order one as they are non-stationary in the levels but stationary in first differences. Hence, the first differences will be employed throughout our econometric analysis. Interestingly, this is in contrast to Lynch and Panigirtzoglou (2003) who reject the non-stationarity hypothesis for the S&P 500 risk-neutral skewness and kurtosis when measured at levels. Yet, results are not comparable. This is because they investigate a different period (1982-2002), use weekly averages rather than weekly data, and employ greater significance levels (10%). In addition, their moments construction algorithm differs from the one employed in this paper.

4. The forecasting models

4.1. The economic variables model

A vast literature documents that asset returns can be predicted using a number of economic variables (see e.g., Goyal and Welch, 2008, and references therein). The results from this stream of research might be exploited for the purpose of forecasting RNMs, as well. This is because the RNMs depend on the mean of the return distribution and thus models that are able to forecast the mean might also be successful in predicting higher moments (see Harvey and Whaley, 1992, and Konstantinidi, Skiadopoulos and, Tzagkaraki, 2008, for an analogous rationale). Therefore, the following economic variables model is employed:

$$\begin{aligned} \Delta MOMENT_{j,t} = c_j + a_j^{+} R_{t-1}^{+} + a_j^{-} R_{t-1}^{-} + \beta_j i_{t-1} + \gamma_j fx_{t-1} + \delta_j oil_{t-1} \\ \varsigma_j \Delta HV_{t-1} + \rho_j \Delta MOMENT_{j,t-1} + \kappa_j \Delta ys_{t-1} + \xi_j vol_{t-1} + \varepsilon_{j,t} \end{aligned} \quad (11)$$

where $\Delta MOMENT_{j,t}$ denotes the daily change from time (t-1) to t in the jth risk-neutral moment j=*MFIV, SKEW, KURT*, c_j is a constant and $R_{t-1}^{+} / R_{t-1}^{-}$ is the positive/negative time t-1 return on the S&P 500 so as to capture the presence of the leverage effect.[11] i_{t-1} denotes the LIBOR rate with time to maturity corresponding to the constant maturity of the moment under consideration, fx_{t-1} denotes the Euro/US Dollar exchange rate, and oil_{t-1} denotes the WTI crude oil price; all three variables are measured in log-differences. Δys_{t-1} denotes the change in the slope of the yield curve measured as the difference between the yield of the ten year US government bond and the LIBOR rate i. vol_{t-1} denotes the S&P 500 volume. The historical volatility HV_{t-1} is calculated as the 30 days moving average standard deviation of returns. Furthermore, the lagged difference of the risk-neutral moment

[11] That is, $R_{t-1}^{+} = \begin{cases} R_{t-1} \text{ if } R_{t-1} > 0 \\ 0 \quad \text{if } R_{t-1} < 0 \end{cases}$ and $R_{t-1}^{-} = \begin{cases} R_{t-1} \text{ if } R_{t-1} < 0 \\ 0 \quad \text{if } R_{t-1} > 0. \end{cases}$

under consideration is incorporated into the regression equation in order to capture the reported autocorrelation effects in Section 3.

4.2 Univariate autoregressive and VAR models

Univariate autoregressive and vector autoregressive (VAR) models are used to study the dynamics of risk-neutral moments (see also Konstantinidi, Skiadopoulos, and Tzagkaraki, 2008, for a similar approach in the setting of implied volatility indices). An AR(1) specification is employed to avoid overfitting the data. The univariate specification is given by

$$\Delta MOMENT_{j,t} = c_j + \lambda_j \Delta MOMENT_{j,t-1} + \varepsilon_{j,t}. \tag{12}$$

Thus, according to this specification, future risk-neutral moments are assumed to be only influenced by the past values of the *same* moment. On the other hand, the VAR specification allows capturing any cross-moments effects in the dynamics of risk-neutral moments; Lynch and Panigirtzoglou (2008) Chang, Christoffersen, and Jacobs (2010), and Hansis, Schlag, and Vilkov (2010) document cross-correlations between risk-neutral moments. Hence, a VAR(1) specification is used to model their dynamics, i.e.

$$\begin{pmatrix} \Delta MFIV_t \\ \Delta SKEW_t \\ \Delta KURT_t \end{pmatrix} = C + \Theta_1 \begin{pmatrix} \Delta MFIV_{t-1} \\ \Delta SKEW_{t-1} \\ \Delta KURT_{t-1} \end{pmatrix} + \varepsilon_t \tag{13}$$

where C is a (3x1) vector of constants and Θ_1 is a (3x3) matrix of autoregressive parameters.

4.3. Cointegration and the vector error correction model

In Section 2, RNMs are found to be integrated of order one. However, a long-run relationship may exist, i.e. they may be cointegrated and consequently there may be a linear combination of RNMs that is stationary in the levels. Such cointegration relations reflect long-rung equilibrium relationships between the variables under scrutiny. In the case where there is cointegration, a vector error correction model (VECM) instead of a VAR model should be used. Therefore, to investigate the presence of any cointegrating relationships in the risk-neutral moments series, a VECM(1) model is employed, i.e.

$$\begin{pmatrix} \Delta MFIV_t \\ \Delta SKEW_t \\ \Delta KURT_t \end{pmatrix} = C + A_1 \begin{pmatrix} \Delta MFIV_{t-1} \\ \Delta SKEW_{t-1} \\ \Delta KURT_{t-1} \end{pmatrix} + \Pi \begin{pmatrix} MFIV_{t-1} \\ SKEW_{t-1} \\ KURT_{t-1} \end{pmatrix} + \varepsilon_t \tag{14}$$

where C is a (3x1) vector of constants, A_1 is a (3x3) matrix of coefficients that controls the short-run dynamics, and Π is a (3x3) matrix of coefficients controlling the speed of convergence to the long-run equilibrium. Π contains information about whether the three moments are cointegrated. In particular, the rank $r(\Pi)$ of the matrix Π equals the number of cointegrating vectors. In the case $r(\Pi)=3$ (i.e. full rank), all moments are integrated of order zero and the appropriate model is a VAR specified in terms of levels. If $r(\Pi)=0,$ then there is no cointegration between the moments and a VAR in first differences is appropriate (see equation (12)). In the case of $0 < r(\Pi) < 3$ there is cointegration and $r(\Pi)=r$ cointegrating vectors exist. In this case, the VECM(1) described by equation (14) should be used.

To determine the number of cointegrating vectors r Johansen's (1988) test for cointegration is employed. In particular, the trace $\lambda_{trace}(r) = -n\sum_{i=r+1}^{K} ln\left(1-\hat{\lambda}_i\right)$ and the maximum $\lambda_{max}(r,r+1) = -n\,ln(1-\hat{\lambda}_{r+1})$ test statistics are used where K is the number of variables in the system, $\hat{\lambda}_i$ is the i^{th} estimated eigenvalue of Π and n is the number of observations. The respective hypotheses are $H_0: r(\Pi) \leq r$ vs. $H_1: r(\Pi) > r$ for the trace statistic and $H_0: r(\Pi) = r$ vs. $H_1: r(\Pi) = r+1$ for the maximum test statistic. Critical values for the test statistics are obtained from Ostrewald-Lenum (1992).

4.4. Univariate ARIMA and ARFIMA models

Univariate ARIMA(p,d,q) and ARFIMA(p,d,q) models are used to investigate short-run dynamics as well as the presence of long-memory effects (see Baillie, 1996, for a review on fractional integration). Konstantinidi, Skiadopoulos, and Tzagkaraki (2008) find that these models are useful to forecast implied volatilities; we extend their analysis to a higher order moments setting. The ARIMA(p,d_j,q) model is described by

$$\Phi_j(L)\Delta^{d_j} MOMENT_{j,t} = c_j + \Theta_j(L)\varepsilon_{j,t} \qquad (15)$$

where d_j denotes the order of differencing that is necessary to produce a stationary and invertible process for the jth moment. In our case, $d_j = 1$ for all j as all moments have been found to be integrated of order one. L denotes the lag operator, $\Phi_j(L) = 1 + \phi_{j,1}L + ... + \phi_{j,p}L^p$ and $\Theta_j(L) = 1 + \theta_{j,1}L + ... + \theta_{j,p}L^p$ the autoregressive and moving average polynomials, respectively, and

$c_j = -\mu_j(1+\phi_{j,1}+...+\phi_{j,p})$ with μ_j being the mean of $\Delta MOMENT_{j,t}$. To avoid overfitting the data we choose $p=q=1$ and hence estimate an ARIMA(1,1,1) model.

The ARFIMA(p,d,q) model is defined as

$$\Phi_j(L)(1-L)^{d_j}(\Delta MOMENT_{j,t}-\mu_j)=\Theta_j(L)\varepsilon_{j,t} \qquad (16)$$

where d_j denotes the order of fractional integration of the jth moment, $(1-L)^{d_j}$ is the fractional difference operator and μ_j is the expected value of $\Delta MOMENT_{j,t}$. If $|d_j|<0.5$ the ARFIMA(p,d,q) model is invertible and second-order stationary.[12] In particular, if $0<d_j<0.5(-0.5<d_j<0)$ the process exhibits a "long memory" (antipersistence) in the sense that the sum of the autocorrelation functions diverges to infinity (a constant). Again, p=q=1 to avoid overfitting the data. Following Pong, Shackleton, and Taylor (2004), maximum likelihood estimation of the ARFIMA$(1,d_j,1)$ model is performed in the frequency domain by using the Whittle approximation of the Gaussian log-likelihood. Based on the estimated model, forecasts are formed by taking the infinite autoregressive expansion of the ARFIMA$(1,d_j,1)$ process. Hence, one step ahead forecasts are obtained by

$$E(MOMENT_{j,t+1}|I_t)=MOMENT_{j,t}+\mu_j-\sum_{k=1}^{\infty}\pi_{k,j}(\Delta MOMENT_{j,t-k+1}-\mu_j) \qquad (17)$$

with $\pi_{k,j}=\sum_{i=0}^{k}(b_{i,j}+\phi_j b_{i-1,j})(-0)^{k-i}$ and $b_i=\dfrac{\Gamma(-d_j+i)}{\Gamma(-d_j)\Gamma(i+1)}$ where $\Gamma(\cdot)$ denotes the gamma function. For the empirical implementation of equation (16) the summation is truncated at 150 as in Konstantinidi, Skiadopoulos, and Tzagkaraki (2008).

5. In-sample evidence

We first estimate all models described in Section 4 using the daily computed RNMs over the in-sample part of the dataset from January 3rd 1996 to January 3rd 2000. Table 3 shows the in-sample estimation results for the economic variables model across all moments and maturities. Newey-West *t*-statistics are reported within brackets. We can see that the economic variables are not statistically significant different from zero in almost all cases. The lagged dependent variable exhibits statistical significance in all cases but the 30-days risk-neutral volatility. Interestingly, even though the adjusted

[12] Note that in contrast to the ARIMA$(1,d_j,1)$, the order of (fractional) integration for the ARFIMA(p,d_j,l) is not predetermined. Instead, it is estimated as part of the estimation of the model. Hence in general, it will not be the same across moments.

R^2 is close to zero for the risk-neutral volatility models, it is substantially greater for the higher moments.

Table 4 shows the estimation results for the univariate autoregressive (Panel A) and VAR model (Panel B). We can see that the univariate autoregressive model does not fit well the in-sample dynamics of the risk-neutral volatility for all maturities under scrutiny. This is in accordance with the findings of Konstantinidi, Skiadopoulos, and Tzagkaraki (2008). On the other hand, the adjusted R^2 increase significantly for the risk-neutral skewness and kurtosis for all maturities ranging from 0.17 to 0.27. Consequently, autoregressive effects seem to be much more important for higher moments than for the risk-neutral volatility. In the case of the VAR model, the adjusted R^2 are greater than the ones obtained in the univariate case. This suggests that there are significant cross-moments effects as also indicated by the fact that there are statistically significant parameters off the main diagonal of Θ_l.

Table 5 shows the estimation results for the VECM(1). Application of Johansen's test yields $r(\Pi)=2$ for all maturities and significance levels; the only exception occurs for the 30-days moments where $r(\Pi)=1$ at the 1% level. Therefore, a VECM(1) with two cointegrating vectors is estimated for each maturity. The adjusted R^2 of the different VECM(1) across the various maturities shows that there is strong in-sample evidence for modelling the dynamics of risk-neutral moments by a VECM(1) specification. Even though the adjusted R^2's for the risk-neutral volatility are lower than these obtained in the VAR setup, they are substantially increased for the higher moments ranging from 0.24 to 0.37.

Table 6 shows the estimation results for the ARIMA*(1,1,1)* model applied to the three risk-neutral moments across the three maturities. We can see that the adjusted R^2*'s* are low for the risk-neutral volatility indicating only weak in-sample evidence. This finding again is in line with the results of Konstantinidi, Skiadopoulos, and Tzagkaraki (2008). However, the in-sample goodness-of-fit of the model is stronger for the risk-neutral skewness and kurtosis as indicated by the adjusted R^2 ranging from 0.10 to 0.19. Especially the MA part of the ARIMA*(1,1,1)* model for risk-neutral skewness and kurtosis is statistically significant in all cases (with the exception of the 30-days measures) while the AR part is only highly statistically significant for the risk-neutral volatility (again with the exception of the 30-days measures).

Finally, Table 7 shows the results for the ARFIMA$(1,d_j,1)$ model. We can see that the fractional differencing parameter d_j is significantly different from zero in only a few cases. Moreover, it lies between 0 and 0.5 only for the 60-days risk-neutral kurtosis indicating that long-memory effects are only present for this single case. Strong statistical significance is found for the moving average polynomial parameter θ in many cases; this coincides with the findings for the MA part of the ARIMA*(1,1,1)* model.

6. Out-of-sample forecasting performance

We assess the out-of-sample forecasting performance of the models described in Section 4 from January 4th 2000 to October 30th 2009. To this end, we estimate each model recursively by employing a constant rolling window of daily computed RNMs.[13] Then, at each point in time, we form one step ahead point forecasts. In line with Goncalves and Guidolin (2006), Konstantinidi, Skiadopoulos, and Tzagkaraki (2008), and Bernales and Guidolin (2010), we evaluate the accuracy of the obtained forecasts by three commonly used measures, i.e. the root mean squared prediction error (RMSE), mean absolute prediction error (MAE), and mean correct prediction of the direction of change (MCP). To evaluate the statistical significance of the obtained figures, first, we examine formally whether the forecasts derived from each one of the considered models outperform the random walk model that is selected as a benchmark. Next, we identify the set of best forecasting models.

6.1. Relative performance against the random walk

We use the modified Diebold-Mariano (MDM, Harvey, Leybourne, and Newbold, 1997) and a ratio test to assess whether any model under consideration outperforms the random walk model in a statistically significant sense under the RMSE/MAE and the MCP metrics, respectively. This is done for each one of the three maturities. MDM tests the null hypothesis of equal predictive accuracy of a given forecasting model and a benchmark model. To fix ideas, let (e_{1t}, e_{2t}) denote the one-step ahead forecasting errors of a given model and a benchmark model, respectively, at any point in time t. These errors are to be evaluated by a specified loss function $g(e)$. Then, the null hypothesis H_0 of equal expected forecasting performance is

$$H_0 = E[g(e_{1t}) - g(e_{2t})] = 0 \qquad (18)$$

Let $d_t = g(e_{1t}) - g(e_{2t})$ and the sample mean $\bar{d} = T^{-1}\sum_{t=1}^{T} d_t$. The variance of $\bar{d}$ is given by $V(\bar{d}) \approx T^{-1}\left[\gamma_0 + 2\sum_{k=1}^{h-1}\gamma_k\right]$ where γ_k denotes the k^{th} autocovariance of d_t and h the forecasting

[13] The first estimation window spans the period January 4th 1996 to January 3rd 2000 and the first forecast refers to January 4th 2000.

horizon. This can be estimated by $\hat{\gamma}_k = T^{-1} \sum_{t=k+1}^{n} (d_t - \bar{d})(d_{t-k} - \bar{d})$. The null hypothesis can be tested by the Diebold-Mariano (1995) test statistic S_1 given by

$$S_1 = \left[V(\bar{d}) \right]^{-1/2} \bar{d} \tag{19}$$

Under the null hypothesis this statistic follows asymptotically a standard normal distribution. However, Harvey, Leybourne, and Newbold (1997) find that the Diebold-Mariano test described by equation (19) can be seriously over-sized in small and moderate samples. Therefore, they propose the following modified version S_1^* of the Diebold-Mariano test statistic

$$S_1^* = \left[\frac{T + 1 - 2h + n^{-1} h(h-1)}{T} \right]^{1/2} S_1 \ . \tag{20}$$

As a further extension, they also suggest the use of the Student's t distribution with (n-1) degrees of freedom instead of the normal distribution for the determination of the critical values for S_1*.

Table 8 reports the values for the three error metrics for each model, moment, and maturity over the chosen out-of-sample period. One and two asterisks denote rejection of the null hypothesis at the 5% and 1% significance levels, respectively. Regarding the predictability of the risk-neutral volatility, we can see that no model outperforms the random walk model under the RMSE and the MAE metrics. However, in the case where the MCP metric is considered, we can see that for each maturity there are models that outperform the random walk in a statistically significant manner. This is in line with the findings of Konstantinidi, Skiadopoulos, and Tzagkaraki (2008) who find that the random walk can be outperformed in forecasting the evolution of the VIX under the MCP metric. In the case of implied skewness and kurtosis, all models but the ARFIMA(1,d_j,1) outperform the random walk under all metrics. The results suggest that there are strong predictable patterns in the evolution of the higher order RNMs.

6.2. Choosing the best models: The model confidence set

The MDM test is constructed to compare the performance of two competing models and hence is not able to select a set of best forecasting models. Instead, the recently developed model confidence set (MCS) methodology by Hansen, Lunde, and Nason (2010) is able to do so. Moreover, the MCS methodology acknowledges the information content in the underlying data in the sense that uninformative data yield a higher number of best models than informative data. To fix ideas, let M_0

denote a set that contains a finite number of forecasting models where each model is indexed by $i = 1,...,m_0$. Each model is associated with a series of forecasting errors denoted by $e_{i,t}$ and evaluated by a loss function $g(\cdot)$. The performance of model *i* relative to *j* is measured at any point in time *t* by the relative performance variable $d_{ij,t} \equiv g(e_{i,t}) - g(e_{j,t})$ for all $i, j \in M_0$. The set M^* of superior models in M_0 is defined by

$$M^* \equiv \left\{ i \in M_0 : E(d_{ij,t}) \leq 0 \;\; for\, all \;\; j \in M_0 \right\}. \tag{21}$$

Hence, choosing the best forecasting model(s) is equivalent to determining M^*. This is done by a sequential series of significance tests that remove one inferior model from the set of candidate models in each testing run. Thus, the MCS procedure starts off with the full set of candidate models M_0 and drops models until it reaches M^*. In order to judge statistically whether M^* has been reached or not the following hypothesis tests have to be carried out in each run

$$H_{0,M} : E(d_{ij,t}) = 0 \;\; for\, all\, i, j \in M \quad vs. \quad H_{A,M} : E(d_{ij,t}) \neq 0 \;\; for\, some\, i, j \in M \tag{22}$$

where $M \subset M_0$ is the set of remaining candidate models for M^*. If $H_{0,M}$ is rejected in an iteration a model has to be removed from the set of candidates M and the hypotheses are tested again on the reduced set of models. This procedure is repeated until $H_{0,M}$ is not rejected at some predetermined confidence level α. The set of models included in M at this point is denoted by $\hat{M}^*_{1-\alpha}$ and referred to as the Model Confidence Set. In order to empirically implement the MCS procedure one needs an *equivalence test* δ_M that allows testing the $H_{0,M}$, as well as an elimination rule e_M that selects the model to remove from *M* if necessary. The equivalence test δ_M is based on the range statistic defined as

$$T_M \equiv \max_{i,j \in M} \left| t_{ij} \right| \tag{23}$$

where $t_{ij} = \dfrac{\overline{d_{ij}}}{\sqrt{var(\overline{d_{ij}})}}$ with $\overline{d_{ij}} = T^{-1} \sum_{t=1}^{T} d_{ij,t}$. The elimination rule associated with T_M is set to be the following: the *i* model selected by e_M to be eliminated is the one for which $t_{e_M,j} = T_M$ for some $j \in M$. The asymptotic distribution of the test statistic T_M under the null hypothesis is non-standard.

Hence, we bootstrap the distribution of the test statistic by implementing the algorithm suggested by Hansen, Lunde, and Nason (2010).[14]

We apply the MCS methodology to all RNMs and maturities using the RMSE and MSE as the error metric, separately. Table 9 reports the models that are included in the model confidence set for the cases of a confidence level α=0.25, 0.1 and 0.05. We can see that in the case of the risk-neutral volatility, the MCS contains the random walk model for every maturity and under both evaluation metrics; in general, the MCSs are quite large containing at least five models. This is consistent with the MDM results where no model outperforms the random walk model.

Regarding the case of implied skewness and kurtosis, the MCS are smaller for every maturity compared to the volatility case. This suggests that some of the examined models forecast risk-neutral skewness and kurtosis significantly better than others. In particular, the largest risk-neutral skewness MCS is encountered in the 60-days maturity case. This contains five models (AR, economic variables, VAR, VECM, ARIMA) in the case where the RMSE is considered, and four models (AR, VAR, VECM, ARIMA) in the case where the MAE is considered. The results hold at any significance level. The findings are similar for the 60-days risk-neutral kurtosis. In particular, the kurtosis MCS consists of the AR, VAR, VECM and ARIMA model for the MAE case for every significance level whereas it additionally contains the economic variables model in the RMSE case (except for the 0.25 level). The fact that both the 60-days risk-neutral skewness and kurtosis MCS are quite large comes at no surprise given that the error metrics in Table 8 are quite similar across models.

Interestingly, the MCS is singleton for the 90-days risk-neutral skewness and kurtosis consisting of the ARIMA*(1,1,1)* model only. Thus, the ARIMA forecasts best the 90-days risk-neutral skewness and kurtosis among all considered models. In fact, the ARIMA model is the only model that is contained in the risk-neutral skewness and kurtosis MCS for every maturity. Hence, we conclude that the ARIMA model is the only model that performs consistently best in forecasting changes in higher order RNMs.

---

[14] In particular, we use a fixed block bootstrap to estimate the distribution of T_M under the null hypothesis. The fixed block bootstrap constructs resamples by concatenating different blocks of length l which are drawn from the original sample with n observations. The number of blocks necessary to construct one resample is given by the smallest integer v satisfying $v > \frac{n}{l}$. The first element of each block in one resample is drawn randomly from the original sample and the remaining $l-1$ elements in a given block are taken to be the $l-1$ consecutive elements from the original sample. Alternatively, another bootstrap scheme such as the stationary bootstrap can be employed. We construct 2,500 bootstrap resamples and the fixed block length is set to l=15. Different choices for l between 2 and 60 have been tested and the composition of the resulting MCS turned out to be robust.

7. Economic significance

7.1. Trading strategies: Design and implementation

The results on the predictability of the risk-neutral moments suggest that the dynamics of risk-neutral skewness and kurtosis can be forecasted under a statistical setting. To explore the economic significance of the obtained forecasts, we develop option trading strategies based on the formed risk-neutral skewness and kurtosis forecasts (also termed skewness and kurtosis strategies). The trading strategies are based on the out-of-sample forecasts derived by each one of models described in Section 4.

We design the proposed strategies being based on the relation of the two higher order RNMs to the shape of the implied volatility curve measured as a function of moneyness on any given day and for any given maturity. In particular, the slope of the implied volatility curve is related to the risk-neutral skewness. The more negative the risk-neutral skewness, the steeper the implied volatility curve is (see also BKM, 2003, for an empirical study on this). Therefore, if the implied skewness is forecasted to decrease this is equivalent to forecasting that the implied volatility skew will steepen and vice versa. Hence, in this case the strategy will be to buy OTM puts and sell OTM calls. On the other hand, the risk-neutral (excess) kurtosis is related to the curvature of the implied volatility curve. So for instance, if implied kurtosis is forecasted to increase, then the strategy would be to sell near-to-the-money options and buy at-the-money and away from the money options.

To fix ideas, let $skew_t$ and $kurt_t$ denote the risk-neutral skewness and kurtosis, respectively, implied by option prices at time t and $\widehat{skew}_{t,t+1}$ and $\widehat{kurt}_{t,t+1}$ denote the point forecast for the risk-neutral skewness and kurtosis, respectively, that will prevail at t+1 as formed at time t. At any point in time t, we form the implied skewness trading strategy based on the following rules: (1) If $skew_t > \widehat{skew}_{t,t+1}$ sell all available OTM calls and buy all available OTM puts, (2) if $skew_t < \widehat{skew}_{t,t+1}$ buy all available OTM calls and sell OTM puts, (3) if $skew_t = \widehat{skew}_{t,t+1}$ do nothing. Similarly, we develop the implied kurtosis trading strategy as follows: (1) If $kurt_t > \widehat{kurt}_{t,t+1}$ buy all available near to-the-money calls and puts and sell all available ATM and deep OTM calls and puts. (2) If $kurt_t < \widehat{kurt}_{t,t+1}$ sell all available near to-the-money calls/puts and buy all available ATM and deep OTM calls and puts. (3) If $kurt_t = \widehat{kurt}_{t,t+1}$ do nothing. We conduct the skewness and kurtosis trading strategies by employing the 60 and 90-days maturity actual and forecasted implied moments; the 30-days maturity actual and forecasted moments are not used since there are less data for these moments and hence using them would result in low trading activity. Also only OTM options are considered since these have greater liquidity than ITM options. To immunise the strategies' returns against changes in the underlying asset price and its volatility, we perform both delta and vega hedging (see Bali and Murray, 2011, for a similar approach).[15] To this end, we vega hedge the option portfolios using either the closest to-the-money short or long option in the portfolio as the hedging instrument depending on whether a short or long vega position is required to set up the hedge. Subsequently, the vega hedged portfolio is delta hedged using the S&P 500 as the hedging instrument in line with Bakshi, Cao, and Chen (1997). Finally, the return on the position is computed as the change in the value of the portfolio (option position and underlying position) from t to t+1 divided by the value at t.

A subtle remark is in order at this point. The extracted RNMs (as well as the point forecasts) refer to synthetic constant maturity options that are not traded in the market. However, the traded options that need to be employed by the strategies have a decreasing time-to-maturity as time goes by. Hence, in principle, actual and forecasted RNMs that match the maturities of the traded options should be employed. To address this issue, we select all traded expirations between 60 and 90 days. In the case where more than one available expirations trade in this interval, the expiration with the greatest number of option quotes is chosen. Next, at any given time t, we perform linear interpolation separately across the 60 and 90 days extracted moments, as well as across the obtained forecasted moments. This enables obtaining actual and forecasted moments that match the chosen maturity of the traded options. Then, we compare these interpolated forecasts to the interpolated risk-neutral moments and set up the trading strategies.[16] To this end, at any point in time t, we define options as being ATM, OTM, far OTM and near OTM by measuring moneyness $m = \frac{K}{S_t}$ and applying the definitions provided in Table 10, where K denotes the strike and S_t the underlying price. For any given moneyness range, all traded options are included in the option portfolio.

[15] Bali and Murray (2011) also implement risk-neutral skewness strategies by constructing "skewness assets ". These are delta and vega hedged positions whose value changes with risk-neutral skewness. The construction principle of their skewness strategies is similar to our trading rules, e.g., in the case where the skewness is forecasted to decrease buy OTM calls and short OTM puts. Their strategies differ from ours though with respect to the number of options involved in the trade. While we trade all available OTM options, they only incorporate (i) one OTM call/put (in case of their CALL/PUT asset) or (ii) one OTM call and one OTM put (in case of their PUTCALL asset). Hence, our skewness strategy nests their skewness assets to a certain extent.

[16] If there is no available maturity between 60 and 90 days no trading takes place. Similarly, if there is no extracted moment at t or no forecast for t+1 no trading takes place either.

7.2. Performance evaluation

We evaluate the performance of the trading strategies by using the Sharpe Ratio (SR) and Leland's alpha (1999, see also Goncalves and Guidolin, 2006, Konstantinidi, Skiadopoulos and Tzagkaraki, 2008, Santa-Clara and Saretto, 2009, for a similar approach), separately. The SR is defined as

$$SR = \frac{E[R_{Strategy} - R_{riskfree}]}{\sqrt{Var(R_{Strategy} - R_{riskfree})}} \tag{24}$$

where $R_{Strategy}$ denotes the return on the trading strategy and $R_{riskfree}$ the risk free rate. The expected value and the variance in equation (26) are replaced by their sample estimates and the risk-free rate is set equal to the US LIBOR rate. Leland's (1999) alpha accounts for the presence of (not reported) non-normalities in the strategy returns. Therefore, it is a more robust measure of alpha compared with Jensen's alpha that uses the CAPM as a model to measure equilibrium returns. To fix ideas, Leland's alpha A_p is defined as

$$A_p = E[R_{Strategy}] - B_p(E[R_{Market}] - R_{riskfree}) - R_{riskfree} \tag{25}$$

$B_p = \frac{Cov(R_{Strategy,} -(1+R_{Market})^{-\gamma})}{Cov(R_{Market,} -(1+R_{Market})^{-\gamma})}$ is a measure of risk similar to the CAPM beta and $\gamma = \frac{ln(E[1+R_{Market}]) - ln(1+R_{riskfree})}{Var[ln(1+R_{Market})]}$ is a measure of risk aversion. We approximate the return on the market portfolio by the S&P 500 return and estimate A_p in a two step procedure. First, we compute γ and B_p at each time step. Second, we estimate the following regression to determine the value of A_p:

$$R_{Strategy,t} - B_{p,t}(R_{Market,t} - R_{riskfree,t}) - R_{riskfree,t} = A_p + \varepsilon_t \tag{26}$$

If $A_p > 0$, then the trading strategy yields an expected return in excess of its equilibrium risk adjusted level.

Table 11 reports the annualised SR and Leland's alpha for the risk-neutral skewness and kurtosis trading strategies over the out-of-sample period from January 2000 to October 2009. Panels A to F report results for the strategies implemented based on the daily forecasts obtained by the respective models described in Section 4; The ARFIMA model is omitted since it did not outperform the random walk model in the comparison under the statistical setting. The SR of the S&P500 buy and hold strategy is also reported. To assess whether the performance measures are statistically different from zero, we provide bootstrapped 95% confidence intervals for the SR and Leland's alpha. We calculate these by employing the Politis and Romano (1994) stationary bootstrap method.[17]

We can see that in the case of trading skewness, all but the random walk model generate economically significant returns the SRs and Leland's alphas are positive and significantly different from zero. This extends the finding reported in Section 6 where these models outperform the random walk under the statistical setting. Moreover, all models but the random walk outperform the S&P 500 buy and hold strategy which yields an SR of -0.83. Interestingly, the ARIMA(1,1,1) forecasts yield the greatest SR of all models. Again, this is in line with the statistical superiority of the ARIMA(1,1,1) model documented in Section 6. However, the SRs and Leland's alphas are quite similar across the other models. This is not surprising given that the values for the MCP metric in Table 8 are also quite similar across these models; the MCP metric is the relevant metric in the context of the skewness and kurtosis strategies as these are based on the forecasted direction of change in the moments. With respect to trading kurtosis, we can see that none of the considered models generates economically significant returns as the performance measures are not significantly positive.

The evidence that trading higher RNMs yields economically significant profits questions the efficiency of the S&P 500 options market. To fully address this, we analyze the profitability of the trading strategies proposed in Section 7.1by incorporating transaction costs. In particular, in line with Bernales and Guidolin (2010), we implement once more the trading strategies using the quoted bid and ask option prices provided by Optionmetrics. Table 12 reports the performance metrics for the two trading strategies once transaction costs are incorporated. We can see that the values for the SR and Leland's alpha are significantly negative for both the skewness and kurtosis trading strategies regardless of the model used to generate the RNM forecasts. Hence, the hypothesis of the efficiency of the S&P 500 options market cannot be rejected.

[17] The confidence intervals for the SR and Leland's alpha are obtained by bootstrapping in order to account for the non-normalities in the strategy returns. Given the low autocorrelation of the strategy returns (not reported here) the average block size is set to ten. Alternative average block sizes were tested and the results proved to be robust against modifications of these block sizes.

8. Further analysis: The effect of the subprime crisis

To check the robustness of the reported findings, we apply the analysis described in Sections 6 and 7 to the subprime crisis period from January 3rd 2007 to October 30th 2009. This check is of interest for the purposes of our analysis given that the behaviour of higher-order RNMs changes over periods of turbulence (see e.g., Lynch and Panigirtzoglou, 2008). Table 13 reports the results for the three error metrics, as well as the results of the MDM test over this sub-sample period. Just as was the case with the full out-of-sample period, the AR(1), VAR(1), and the ARIMA(1,1,1) models outperform the random walk in forecasting the risk-neutral skewness and kurtosis for each maturity. The economic variables and VECM model outperform the random walk in forecasting the risk-neutral volatility for all maturities.

Table 14 reports the composition of the MCS over the subprime crisis sub-sample across moments, maturities, and significance levels. We can see that the MCS tend to be larger for the crisis sample compared to the full sample case results reported in Table 9. In particular, in contrast to the full sample case, the 90-days risk-neutral skewness and kurtosis MCS are not any longer singleton for all significance levels. This is because the MCS methodology acknowledges the informational content of the data: given that a smaller dataset is used, the data contain less information and consequently a larger MCS results. Nevertheless, the risk-neutral skewness and kurtosis MCS are still singleton at the 25% level containing only the ARIMA*(1,1,1)* model. Unreported results on the performance of the skewness and trading strategies show that the SR and Leland’s alpha are insignificant for both strategies. Hence, despite the statistical predictability, the two trading strategies do not yield economically significant profits during the crisis. This is not surprising given that the market was too volatile over the crisis period.

9. Summary and conclusions

We investigate whether the dynamics of higher order risk-neutral moments (RNMs) can be used to forecast their subsequent movements. This provides an alternative way of exploring whether the evolution of implied volatility surfaces (IVSs) over time can be forecasted and contributes to the related literature. To this end, we compute the S&P 500 RNMs by the model-free method of Bakshi, Kapadia and Madan (2003) and then perform a horse race among alternative model specifications to investigate their statistical and economic *out-of-sample* forecasting performance. To conduct our study, we extract 30, 60 and 90-days constant maturity risk-neutral volatility, skewness and kurtosis from the S&P 500 market option prices over the period 1996 to 2009. To the best of our knowledge, this is the first study that addresses whether higher order RNMs are predictable per se by undertaking this research approach.

To identify the process that governs the evolution of the RNMs, we study the out-of-sample forecasting performance of a number of models. In particular, we employ six alternative models (economic variables, AR(1), VAR(1), VECM(1), ARIMA*(1,1,1)* and ARFIMA*(1,d,1)*) in the forecasting horse race. We evaluate the forecasting performance of each one of these models under a number of statistical measures by using formal statistical tests. Moreover, we select the best performing models by employing the newly developed Model Confidence Set (MCS) methodology of Hansen, Lunde and Nason (2010). Furthermore, we assess the economic significance of the obtained higher order RNMs forecasts by means of skewness and kurtosis option trading strategies.

We find that higher RNMs can be forecasted statistically while this is not the case for risk-neutral volatility in general. In particular, all models but the ARFIMA*(1,d,1)* significantly outperform the random walk in forecasting the evolution of higher RNMs. In addition, the application of the MCS test shows that the ARIMA*(1,1,1)* model is consistently among the best performing models across the various maturities. The skewness trading strategy yields economically significant profits. However, the economic significance vanishes once we incorporate transaction costs. Interestingly, the statistical significance of the formed forecasts prevails over the subprime crisis period.

Our findings have at least four implications. First, the reported predictable patterns in the dynamics of higher order RNMs imply that the evolution of implied volatility surfaces is also predictable hence extending the results of prior related literature. Second, the past values of higher-order RNMs affect their dynamics over time. This supports the approach taken by the existing literature on modelling the dynamics of the physical higher moments where autoregressive patterns in their evolution are allowed (e.g., Hansen, 1994, Harvey and Siddique, 1999, Jondeau and Rockinger 2003, Brooks, Burke, Heravi, and Persand, 2005). Third, the reported findings suggest that cross-moments effects and cointegration among moments should also be taken into account for the purposes of modelling higher moments. This can also be proven to be helpful for the development of option pricing models that will be able to respect the evolution of IVSs over time. Fourth, in principle, profitable strategies based on trading skewness may be developed. However, the presence of bid-ask spreads eliminates any abnormal profits confirming the efficiency of the S&P 500 option market. This is in line with the findings in Gonçalves and Guidolin (2006) and Santa-Clara and Saretto (2009).

This paper has taken a comprehensive analysis of the dynamics of higher RNMs. A possible extension would be the analysis of option markets written on other underlying indexes and individual equities. Moreover, the set of forecasting models employed in this work might be augmented by other specifications, as well. In the interest of brevity, these topics are best left for future research.

References

Ait-Sahalia, Y., Lo, A.W., 1998. Nonparametric estimation of state-price densities implicit in financial asset prices. Journal of Finance 53, 499-547.

Ait-Sahalia, Y., Wang, Y., Yared, F., 2001. Do option markets correctly price the probabilities of movement of the underlying asset? Journal of Econometrics 102, 67-110.

Alexander, C., Lazar, E., 2006. Normal mixture GARCH(1,1): Applications to foreign exchange markets. Journal of Applied Econometrics 21, 307-336.

Alexander, C., 2008. Market risk analysis, Volume III: Pricing, hedging and trading financial instruments. John Wiley & Sons.

Ang, A, Bali, T.G., Cakici, N., 2010. The joint cross section of stocks and options. Working paper, Baruch College.

Baillie, R. T., 1996. Long memory processes and fractional integration in econometrics. Journal of Econometrics 73, 5-59.

Bakshi, G., Cao, C. and Chen, Z. 1997. Empirical performance of alternative option pricing models. Journal of Finance 52, 2003-2049.

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

Bakshi, G., Panayotov, G., Skoulakis, G., 2010. Improving the predictability of real economic activity and asset returns with forward variances inferred from option portfolios. Journal of Financial Economics, forthcoming

Bali, T.G., and Hovakimian, A., 2009. Volatility spreads and expected stock returns. Management Science 55, 1797-1812.

Bali, T.G., Murray, S., 2011. Does risk-neutral skewness predict the cross-section of equity portfolio returns? Working paper, Baruch College.

Bates, D.S., 2003. Empirical option pricing: A retrospection. Journal of Econometrics 116, 387-404.

Bernales, A., Guidolin, M., 2010. Can we forecast the implied volatility surface dynamics for CBOE equity options? Predictability and economic value tests, Working paper, Manchester University.

Birru, J., Figlewski, S., 2010. Anatomy of a meltdown: The risk neutral density for the S&P 500 in the fall of 2008. Working paper, New York University.

Bliss, R. R., Panigirtzoglou, N., 2002. Testing the stability of implied probability density functions. Journal of Banking and Finance 26, 381-442.

Bliss, R. R., Panigirtzoglou, N., 2004. Option-implied risk aversion estimates. Journal of Finance 58, 407-446.

Breeden, D., Litzenberger, R., 1978. Prices of state-contingent claims implicit in options prices. Journal of Business 51, 621-651.

Broadie, M., Chernov, M., Johannes, M. (2007). Model specification and risk premia: Evidence from futures options. Journal of Finance 62, 1453–1490.

Brooks, C., Burke, S.P., Heravi, S., Persand, G., 2005. Autoregressive conditional kurtosis. Journal of Financial Econometrics 3, 399-421.

Canina, L., Figlewski, S., 1993. The informational content of implied volatility. Review of Financial Studies 6, 659-681.

Carr, P., Wu, L., 2003. The finite moment log stable process and option pricing. Journal of Finance 58, 753-778.

Carverhill, A., Cheuk, T.H.F., Dyrting, S., 2009. The smirk in the S&P 500 futures options prices: A linearized factor analysis. Review of Derivatives Research 12, 109-139.

Chalamandaris, G., Tsekrekos, A.E., 2010. Predictable dynamics in implied volatility surfaces from OTC currency options. Journal of Banking and Finance 34, 1175-1188.

Chang, B.Y., Christoffersen, P., Jacobs, K., Vainberg, G., 2009. Option-implied measures of equity risk. Working paper, McGill University.

Chang, B.Y., Christoffersen, P., Jacobs, K., 2010. Market volatility, skewness, and kurtosis risks and the cross-section of stock returns, Working Paper, McGill University.

Christoffersen, P.F., Heston, S., and Jacobs, K., 2009. The shape and term structure of the index option smirk: Why multifactor stochastic volatility models work so well. Management Science 55, 1914-1932.

Christoffersen, P.F., Jacobs, K., Ornthanalai, C., 2010. Exploring time-varying jump intensities: Evidence from S&P500 returns and options. Working paper, University of Toronto.

Cremers, M. and D. Weinbaum, 2010. Deviations from put-call parity and stock return predictability, Journal of Financial and Quantitative Analysis 45, 335-367.

DeMiguel, V., Y. Plyakha, R. Uppal, G. Vilkov. 2011. Improving portfolio selection using option-implied volatility and skewness. Working paper, Goethe University, Frankfurt am main, Germany

Dennis, P., Mayhew, D., 2002. Risk-neutral skewness: Evidence from stock options. Journal of Financial and Quantitative Analysis 37, 471-492.

Diebold, F. X., Mariano, R. S., 1995. Comparing predictive accuracy. Journal of Business and Economic Statistics 13, 253-263.

Dumas, B., Fleming, J., Whaley, R.E., 1998. Implied volatility functions: Empirical tests. Journal of Finance 53, 2016–2059.

Engle, R., Mistry, A., 2008. Priced risk and asymmetric volatility in the cross-section of skewness. Working Paper, New York University.

Fengler, M.R., Härdle, W.K., Villa, C., 2003. The dynamics of implied volatilities: A common principal components approach. Review of Derivatives Research 6, 179-202.

Fengler, M.R., Härdle, W.K., Mammen, E., 2007. A semiparametric factor model for implied volatility Surface dynamics. Journal of Financial Econometrics 5, 189-218.

Figlewski, S., 2010. Estimating the implied risk neutral density for the U.S. market portfolio. In Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle (Bollerslev, Russell and Watson, eds.). Oxford, UK: Oxford University Press.

Giamouridis, D., and G. Skiadopoulos, 2010. The informational content of financial options for quantitative asset management: A review. In B. Scherer and K. Winston, eds: Handbook of Quantitative Asset Management, (Oxford University Press), forthcoming.

Gonçalves, S., Guidolin, M., 2006. Predictable dynamics in the S&P 500 index options implied volatility surface. Journal of Business 79, 1591-1635.

Goyal, A., Welch, I., 2008. A comprehensive look at the empirical performance of equity premium prediction. Review of Financial Studies 21, 1455-1508.

Goyal, A., Saretto, A., 2009. Cross-section of option returns and volatility. Journal of Financial Economics 94, 310-326.

Haas, M., Mittnik, S., Paolella M.S., 2004. Mixed normal conditional heteroskedasticity. Journal of Financial Econometrics 2, 211–250.

Hansen, B. E., 1994. Autoregressive conditional density estimation. International Economic Review 35, 705-730.

Hansen, P. R., Lunde, A., Nason, J. M., 2010. The model confidence set. Econometrica, forthcoming.

Hansis, A., Schlag, C., Vilkov, G., 2010. The dynamics of risk-neutral moments: Evidence from individual options. Working paper, Goethe University Frankfurt.

Harvey, C. R., Whaley, R. E., 1992. Market volatility prediction and the efficiency of the S&P 100 index option market. Journal of Financial Economics 31, 43-73.

Harvey, C. R., Siddique, A., 1999. Autoregressive conditional skewness. Journal of Financial and Quantitative Analysis 34, 465-487.

Harvey, D., Leybourne, S., Newbold, P., 1997. Testing the equality of prediction mean squared errors. International Journal of Forecasting 13, 281-291.

Jackwerth, J. C., Rubinstein, M., 1996. Recovering probability distributions from option prices. Journal of Finance 51, 1611-1631.

Jackwerth, J., 2004. Option-implied risk-neutral distributions and risk aversion. Research Foundation of CFA Institute.

Jha R., Kalimipalli, M., 2010. The economic significance of conditional skewness forecasts in index option markets. Journal of Futures Markets 30, 378-406.

Jiang, G. J., Tian Y.S., 2005. The model-free implied volatility and its information content. Review of Financial Studies, 1305-1342.

Jiang, G. and Y. Tian, 2007. Extracting model-free volatility from option prices: An examination of the VIX index. Journal of Derivatives 14, 1-26.

Johansen, S., 1988. Statistical analysis of cointegrating vectors. Journal of Economic Dynamics and Control 12, 231-254.

Jondeau, E., Rockinger, M., 2003. Conditional volatility, skewness, and kurtosis: existence, persistence, and comovements. Journal of Economic Dynamics and Control 27, 1699-1737.

Jondeau, E., Rockinger, M., 2008. The economic value of distributional timing. Working paper, Swiss Finance Institute.

Jondeau, E., Rockinger, M., 2009. The impact of shocks on higher moments. Journal of Financial Econometrics 7, 77-105.

Kim, T-H., Manganelli, S., White, H., 2008. Modeling autoregressive conditional skewness and kurtosis with multi-quantile CAViaR. Working paper, European Central Bank.

Konstantinidi, E., Skiadopoulos, G., Tzagkaraki, E., 2008. Can the evolution of implied volatility be forecasted? evidence from european and US implied volatility indices. Journal of Banking and Finance 32, 2401-2411.

Kostakis, A., Panigirtzoglou, N., Skiadopoulos, G., 2011. Market timing with option-implied distributions: A forward-looking approach. Management Science, forthcoming.

Leland, H. E., 1999. Beyond mean-variance: Performance measurement in a nonsymmetrical world. Financial Analysts Journal 55, 27-36.

Lynch, D., Panigirtzoglou, N., 2008. Summary statistics of option-implied probability density functions and their properties. Working paper, Bank of England.

Malz, A. M., 1997. Estimating the probability distribution of the future exchange rate. Journal of Derivatives 5, 18-36.

Merton, R. C. 1973. Theory of rational option pricing. Bell Journal of Economics and Management Science 4, 141-183.

Ostrewald-Lenum, M.. 1992. A note with quantiles of the asymptotic distribution of the maximum likelihood cointegration rank test statistics. Oxford Bulletin of Economics and Statistics 54, 461-471.

Panigirtzoglou, N., Skiadopoulos, G., 2004. A new approach to modeling the dynamics of implied distributions: Theory and evidence from the S&P 500 options. Journal of Banking and Finance 28, 1499-1520.

Patton, A. J., 2004. On the out-of-sample importance of skewness and asymmetric dependence for asset allocation. Journal of Financial Econometrics 2, 130-168.

Politis, D. N., Romano, J. P., 1994. The stationary bootstrap. Journal of the American Statistical Association 89, 1303-1313.

Pong, S., Shackleton, M. B., Taylor, S. J., Xu, X., 2004. Forecasting currency volatility: A comparison of implied volatilities and AR(FI)MA models. Journal of Banking and Finance 28, 2541-2563.

Rehman, Z., Vilkov, G., 2010. Risk-neutral skewness: Return predictability and its sources. Working paper, Goethe University.

Santa-Clara, P., Saretto, A., 2009. Option strategies: Good deals and margin calls. Journal of Financial Markets 12, 391-417.

Skiadopoulos, G., Hodges, S.D., Clewlow, L., 1999. The dynamics of the S&P 500 implied volatility surface, Review of Derivatives Research 3, 263—282.

Taylor, S. J., Yadav, P. K., Zhang, Y., 2009. Cross-sectional analysis of risk-neutral skewness. Journal of Derivatives 16, 38-52.

Xing, Y., X. Zhang, and R. Zhao, 2010. What does individual option volatility smirk tell us about future equity returns?, Journal of Financial and Quantitative Analysis 45, 641-662.

Zhang, J. E., Xiang, Y., 2008. The implied volatility smirk. Quantitative Finance 8, 263-284.

Tables

Moment	MFIV	SKEW	KURT
Panel A: Summary statistics for the 30-days S&P 500 moments			
NA percentage	35.52 %	35.52 %	35.52 %
Mean	0.21	-0.74	3.32
Median	0.2	-0.73	3.30
Standard Deviation	0.09	0.16	0.16
Skewness	1.79	-0.07	0.48
Kurtosis	8.86	3.12	3.01
Panel B: Summary statistics for the 60-days S&P 500 moments			
NA percentage	10.14 %	10.14 %	10.14 %
Mean	0.21	-0.88	3.47
Median	0.20	-0.87	3.44
Standard Deviation	0.08	0.16	0.22
Skewness	1.71	-0.085	0.60
Kurtosis	8.23	2.82	2.92
Panel C: Summary statistics for the 90-days S&P 500 moments			
NA percentage	9.82 %	9.82 %	9.82 %
Mean	0.22	-0.95	3.58
Median	0.21	-0.94	3.53
Standard Deviation	0.08	0.17	0.28
Skewness	1.49	-0.16	0.75
Kurtosis	6.77	2.76	3.13

Table 1 Entries report the summary statistics of the extracted risk-neutral S&P 500 moments over the sample period from January 4th 1996 to October 30th 2009.

	30 Days			60 Days			90 Days		
	MFIV	**SKEW**	**KURT**	**MFIV**	**SKEW**	**KURT**	**MFIV**	**SKEW**	**KURT**
Panel A: Statistical and Time-Series Properties of the risk-neutral moments levels									
# Observations	616	616	616	857	857	857	789	789	789
Mean	0.21	-0.80	3.38	0.21	-0.91	3.51	0.22	-0.97	3.60
Median	0.20	-0.81	3.39	0.21	-0.93	3.50	0.22	-0.99	3.57
Max	0.46	-0.25	3.83	0.44	-0.31	4.21	0.43	-0.37	4.64
Min	0.11	-1.18	2.96	0.11	-1.33	3.02	0.11	-1.50	3.02
Standard Dev.	0.05	0.16	0.16	0.05	0.16	0.18	0.06	0.16	0.23
Skewness	1.48	0.53	-0.13	1.07	0.75	0.12	0.76	0.70	0.80
Kurtosis	6.71	3.36	2.78	4.95	3.55	3.11	4.06	3.74	4.60
ρ_1	0.68**	0.44**	0.39**	0.88**	0.68**	0.65**	0.76**	0.61**	0.58**
ADF	-1.61	-1.45	-5.34**	-0.78	-0.78	-1.62	0.08	-0.66	-1.94
Panel B: Statistical and Time-Series Properties of the risk-neutral moments first differences									
# Observations	448	448	448	751	751	751	651	651	651
Mean	0.0012	0.0025	-0.0008	-0.0002	0.0031	-0.0043	-0.0001	0.0007	-0.0016
Median	0.0006	0.0012	0.0035	-0.0002	0.0030	-0.0047	0.0003	-0.0014	0.0014
Max	0.11	0.41	0.56	0.06	0.42	0.52	0.05	0.34	0.67
Min	-0.07	-0.40	-0.42	-0.06	-0.40	-0.57	-0.04	-0.41	-0.60
Standard Dev.	0.02	0.13	0.15	0.01	0.10	0.13	0.01	0.10	0.15
Skewness	0.89	-0.04	0.18	0.19	0.23	-0.27	0.13	0.03	-0.14
Kurtosis	10.55	3.34	3.68	7.00	4.62	6.05	4.95	4.12	5.23
ρ_1	-0.07	-0.34**	-0.36**	-0.07	-0.37**	-0.37**	-0.08	-0.35**	-0.33**
ADF	-3.5*	-7.5**	-7.5**	-16.8**	-12.0**	-11.8**	-5.2**	-15.3**	-15.0**

Table 2 Entries report summary and time series statistics of the extracted S&P 500 risk-neutral moments in levels and first differences over the in-sample period from January 4th 1996 to January 3rd 2000. ρ_1 and ADF denote the first-order autocorrelation and value of the augmented Dickey Fuller Test statistic, respectively. One and two asterisks denote rejection of the null hypotheses of zero first-order autocorrelation and existence of a unit root at the 5% and 1% significance level, respectively.

Dep. Variable	30 days $\Delta MFIV_t$	30 days $\Delta SKEW_t$	30 days $\Delta KURT_t$	60 days $\Delta MFIV_t$	60 days $\Delta SKEW_t$	60 days $\Delta KURT_t$	90 days $\Delta MFIV_t$	90 days $\Delta SKEW_t$	90 days $\Delta KURT_t$
Incl. Obs.	326	326	326	668	668	668	553	553	553
c1	0.002*	0.012	-0.017	0.000	0.006	-0.011*	0.000	-0.007	0.004
	[2.12]	[1.30]	[-1.92]	[-0.14]	[1.36]	[-2.14]	[0.17]	[-1.23]	[0.53]
R^+_{t-1}	-0.07	-1.01	1.84	0.02	-0.38	1.17*	-0.06	0.64	-0.39
	[-0.46]	[-0.73]	[1.65]	[0.33]	[-0.69]	[2.02]	[-0.80]	[1.06]	[-0.53]
R^-_{t-1}	0.06	0.92	-1.26	0.04	0.46	-0.79	-0.01	-0.62	0.39
	[0.23]	[1.19]	[-1.60]	[0.32]	[1.10]	[-1.32]	[-0.06]	[-0.86]	[0.39]
i_{t-1}	0.40*	-0.48	-0.27	-0.08	0.43	-0.20	0.02	0.20	-0.09
	[1.99]	[-0.3]	[-0.17]	[-1.49]	[1.83]	[-0.71]	[0.26]	[0.28]	[-0.10]
fx_{t-1}	-0.01	1.02	-1.19	-0.01	-0.36	0.43	-0.06	-1.15	1.47
	[-0.08]	[0.83]	[-0.95]	[-0.09]	[-0.54]	[0.55]	[-0.64]	[-1.72]	[1.72]
oil_{t-1}	0.05	-0.44	0.29	0.03	-0.24	0.18	0.04*	-0.28	0.30
	[1.21]	[-1.61]	[0.92]	[1.31]	[-1.61]	[0.95]	[2.22]	[-1.71]	[1.25]
ΔHV_{t-1}	0.36	-0.18	-0.31	0.12	0.42	-0.41	0.09	0.10	-0.02
	[1.42]	[-0.19]	[-0.32]	[0.95]	[1.07]	[-0.83]	[0.87]	[0.25]	[-0.04]
Dep. Variable (t-1)	-0.13	-0.49**	-0.50**	-0.08*	-0.44**	-0.45**	-0.11**	-0.43**	-0.41**
	[-1.62]	[-10.40]	[-10.89]	[-2.03]	[-13.48]	[-12.95]	[-2.78]	[-12.49]	[-11.55]
ys_{t-1}	0.01	-0.09	0.04	0.01	-0.03	0.03	0.02	-0.03	0.01
	[0.69]	[-0.90]	[0.34]	[0.98]	[-0.54]	[0.50]	[1.76]	[-0.56]	[0.11]
vol_{t-1}	-0.01	-0.06	0.07	-0.003	-0.02	0.02	-0.003	-0.01	0.01
	[-1.33]	[-1.22]	[1.280]	[-1.47]	[-0.98]	[0.89]	[-1.52]	[-0.66]	[0.77]
Adj. R^2	0.04	0.24	0.27	0.01	0.19	0.20	0.03	0.18	0.17

Table 3 Entries report the estimation results for the economic variables model applied to each one of the risk-neutral moment series (ΔMFIV, ΔSKEW, ΔKURT) measured in first differences. 30, 60 and 90-days constant maturity moments are considered. Dep. Variable(*t*-1) denotes the first lagged moment measured in first differences. Newey-West *t*-statistics are reported in brackets. One and two asterisks denote rejection of a zero coefficient at the 5% and 1% significance level, respectively. The sample spans the period from January 4th 1996 to January 3rd 2000.

Dep. Variable	30 days $\Delta MFIV_t$	30 days $\Delta SKEW_t$	30 days $\Delta KURT_t$	60 days $\Delta MFIV_t$	60 days $\Delta SKEW_t$	60 days $\Delta KURT_t$	90 days $\Delta MFIV_t$	90 days $\Delta SKEW_t$	90 days $\Delta KURT_t$
Panel A: Univariate Autoregressive models									
Incl. Obs.	326	326	326	668	668	668	553	553	553
c1	0.002*	0.004	-0.005	0.000	0.003	-0.003	0.000	-0.001	0.001
	[2.29]	[0.70]	[-0.71]	[-0.23]	[0.76]	[-0.80]	[-0.06]	[-0.27]	[0.10]
Dep. Variable (*t*-1)	-0.11	-0.52**	-0.52**	-0.08*	-0.44**	-0.46**	-0.01*	-0.43**	-0.41**
	[-1.88]	[-10.30]	[-10.99]	[-1.96]	[-12.67]	[-13.18]	[-2.28]	[-10.96]	[-10.77]
Adj. R^2	0.01	0.24	0.27	0.00	0.19	0.21	0.01	0.18	0.17
Panel B: VAR model									
$\Delta MFIV_{t-1}$	-0.20**	-2.63**	2.22**	-0.08	0.84**	-1.32**	-0.12**	0.36	-0.92
	[-3.16]	[-5.74]	[4.18]	[-1.93]	[2.77]	[-3.54]	[-2.68]	[0.91]	[-1.67]
$\Delta SKEW_{t-1}$	-0.02	-0.32	-0.19	-0.02	-0.31**	-0.10**	-0.03*	-0.48**	0.12
	[-0.88]	[-1.95]	[-1.00]	[-1.22]	[-3.28]	[-0.87]	[-2.13]	[-4.31]	[0.75]
$\Delta KURT_{t-1}$	-0.01	0.13	-0.65**	-0.02*	0.09	-0.51**	-0.02*	-0.05	-0.33**
	[-0.41]	[0.90]	[-4.00]	[-1.95]	[1.22]	[-5.58]	[-2.22]	[-0.59]	[-3.05]
C	0.002*	0.01	-0.01	0.00	0.00	-0.00	0.00	-0.00	0.00
	[2.13]	[0.74]	[-1.00]	[-0.33]	[0.83]	[-0.88]	[-0.07]	[-0.26]	[0.10]
Adj. R^2	0.04	0.30	0.28	0.01	0.20	0.22	0.01	0.18	0.18

Table 4 Entries report the estimation results of the AR(1) model (Panel A) and VAR(1) model (Panel B) applied to each one of the risk-neutral moment series (ΔMFIV, ΔSKEW, ΔKURT) measured in first differences. 30, 60 and 90-days constant maturity moments are considered. Dep. Variable(*t*-1) denotes the first lagged moment measured in first differences. Newey-West *t*-statistics are reported in brackets. One and two asterisks denote rejection of a zero coefficient at the 5% and 1% significance level, respectively. The sample spans the period from January 4th 1996 to January 3rd 2000.

	30 days			60 days			90 days		
Dep. Variable	$\Delta MFIV_t$	$\Delta SKEW_t$	$\Delta KURT_t$	$\Delta MFIV_t$	$\Delta SKEW_t$	$\Delta KURT_t$	$\Delta MFIV_t$	$\Delta SKEW_t$	$\Delta KURT_t$
$MFIV_{t-1}$	-0.02	0.57**	-0.92**	-0.001	0.02	-0.37**	-0.01	0.29**	-0.82**
	[-1.86]	[6.55]	[-8.62]	[-0.08]	[0.23]	[-3.32]	[-1.01]	[2.79]	[-5.64]
$SKEW_{t-1}$	-0.03**	0.39**	-0.76**	-0.001	-0.04	-0.25**	-0.01	0.09	-0.41**
	[-3.14]	[5.44]	[-8.62]	[-0.12]	[-0.65]	[-3.18]	[-1.00]	[1.59]	[-4.97]
$KURT_{t-1}$	-0.03**	0.59**	-1.00**	-0.001	0.089	-0.37**	-0.01	0.21**	-0.51**
	[-3.09]	[8.27]	[-11.57]	[-0.09]	[1.62]	[-5.52]	[-1.36]	[4.65]	[-8.04]
$\Delta MFIV_{t-1}$	-0.20**	-2.84**	2.81**	-0.08	0.78**	-1.06**	-0.11*	0.24	-0.56
	[-3.02]	[-6.08]	[5.29]	[-1.88]	[2.60]	[-2.90]	[-2.52]	[0.64]	[-1.06]
$\Delta SKEW_{t-1}$	-0.03	-0.43*	0.08	-0.02	-0.27**	0.01	-0.02	-0.55**	0.36*
	[-1.07]	[-2.56]	[0.41]	[-1.07]	[-2.68]	[0.04]	[-1.70]	[-4.73]	[2.21]
$\Delta KURT_{t-1}$	-0.01	-0.08	-0.25	-0.02	0.06	-0.34**	-0.02	-0.16*	-0.06
	[-0.66]	[-0.56]	[-1.48]	[-1.71]	[0.73]	[-3.43]	[-1.62]	[-2.01]	[-0.50]
C	0.08**	-1.78**	2.95**	0.00	0.00	0.00	0.02	-0.73**	1.61**
	[3.02]	[-8.99]	[12.25]	[-0.02]	[0.02]	[0.00]	[1.44]	[-5.61]	[8.86]
Adj. R^2	0.02	0.36	0.37	0.01	0.24	0.27	0.01	0.24	0.26

Table 5 Entries report the estimation results for the VECM(1) model applied to each one of the risk-neutral moment series (ΔMFIV, ΔSKEW, ΔKURT) measured in first differences. 30, 60 and 90-days constant maturity moments are considered. Newey-West *t*-statistics are reported in brackets. One and two asterisks denote rejection of a zero coefficient at the 5% and 1% significance level, respectively. The sample spans the period from January 4th 1996 to January 3rd 2000.

	30 days			60 days			90 days		
Dep. Variable	$\Delta MFIV_t$	$\Delta SKEW_t$	$\Delta KURT_t$	$\Delta MFIV_t$	$\Delta SKEW_t$	$\Delta KURT_t$	$\Delta MFIV_t$	$\Delta SKEW_t$	$\Delta KURT_t$
Incl. Obs.	448	448	448	751	751	751	651	651	651
c	0.00	0.00	-0.00	0.00	0.003*	-0.004*	0.00	0.00	-0.00
	[-0.01]	[0.00]	[-0.17]	[-0.57]	[2.26]	[-2.45]	[-0.21]	[0.43]	[-0.59]
φ	-0.00	-0.22	-0.30	0.51**	0.15*	0.13	0.51**	0.07	0.01
	[-0.07]	[-1.23]	[-1.76]	[4.39]	[2.09]	[1.76]	[3.16]	[1.00]	[0.93]
θ	-0.05	-0.13	-0.05	-0.62**	-0.63**	-0.62**	-0.61**	-0.54**	-0.50**
	[1.50]	[-0.72]	[-0.29]	[-6.01]	[-11.25]	[-10.60]	[-4.12]	[-9.34]	[-8.07]
Adj. R^2	0.00	0.11	0.11	0.01	0.19	0.19	0.01	0.17	0.15

Table 6 Entries report the estimation results of the ARIMA(1,1,1) model applied to each one of the risk-neutral moment series (ΔMFIV, ΔSKEW, ΔKURT) measured in first differences. 30, 60 and 90-days constant maturity moments are considered. Newey-West *t*-statistics are reported in brackets. One and two asterisks denote rejection of a zero coefficient at the 5% and 1% significance level, respectively. The sample spans the period from January 4th 1996 to January 3rd 2000.

	30 days			60 days			90 days		
Dep. Variable	$\Delta MFIV_t$	$\Delta SKEW_t$	$\Delta KURT_t$	$\Delta MFIV_t$	$\Delta SKEW_t$	$\Delta KURT_t$	$\Delta MFIV_t$	$\Delta SKEW_t$	$\Delta KURT_t$
Incl. Obs.	448	448	448	751	751	751	651	651	651
d	0.37	-0.12	-0.17**	-0.08*	0.40	0.47**	-0.09*	0.10	0.06
	[1.57]	[-1.89]	[-2.82]	[-2.26]	[1.70]	[4.02]	[-2.20]	[0.812]	[0.56]
φ	-0.30*	0.42*	0.54**	0.22	0.01	0.07	0.34	-0.05	-0.06
	[-2.27]	[2.05]	[3.14]	[0.26]	[0.09]	[0.80]	[0.63]	[-0.74]	[-0.86]
θ	-0.72**	0.20	0.35	0.23	-0.87**	-0.90**	0.36	-0.61**	-0.55**
	[-5.16]	[0.77]	[1.62]	[0.27]	[-9.04]	[-21.43]	[0.69]	[-6.05]	[-5.08]

Table 7 Entries report the estimation results of the ARFIMA(1,*d*,1) model applied to each one of the risk-neutral moment series (ΔMFIV, ΔSKEW, ΔKURT) measured in first differences. 30, 60 and 90-days constant maturity moments are considered. Newey-West *t*-statistics are reported in brackets. One and two asterisks denote rejection of a zero coefficient at the 5% and 1% significance level, respectively. The sample spans the period from January 4th 1996 to January 3rd 2000.

		30 days			60 days			90 days		
		$MFIV_t$	$SKEW_t$	$KURT_t$	$MFIV_t$	$SKEW_t$	$KURT_t$	$MFIV_t$	$SKEW_t$	$KURT_t$
Random Walk	RMSE	0.016	0.11	0.115	0.012	0.077	0.097	0.011	0.072	0.107
	MAE	0.011	0.082	0.082	0.008	0.056	0.068	0.007	0.051	0.072
	MCP	0.500	0.500	0.500	0.500	0.500	0.500	0.500	0.500	0.500
Economic Variables	RMSE	0.016	0.097**	0.101**	0.012	0.068**	0.088**	0.011	0.065**	0.096**
	MAE	0.012	0.072**	0.072**	0.008	0.050**	0.062**	0.007	0.046**	0.066**
	MCP	0.539*	0.636**	0.627**	0.528*	0.642**	0.634**	0.530**	0.644**	0.634**
AR(1)	RMSE	0.016	0.095**	0.099**	0.012	0.068**	0.087**	0.011	0.064**	0.095**
	MAE	0.011	0.070**	0.070**	0.008	0.049**	0.061**	0.007	0.046**	0.065**
	MCP	0.537*	0.646**	0.656**	0.517	0.653**	0.654**	0.534**	0.643**	0.640**
VAR(1)	RMSE	0.017	0.095**	0.099**	0.012	0.068**	0.087**	0.011	0.065**	0.096**
	MAE	0.012	0.071**	0.070**	0.008	0.050**	0.061**	0.007	0.046**	0.065**
	MCP	0.521	0.645**	0.655**	0.527*	0.653**	0.651**	0.540**	0.633**	0.630**
VECM(1)	RMSE	0.0164	0.093**	0.096**	0.012	0.068**	0.087**	0.011	0.065**	0.096**
	MAE	0.011	0.070**	0.070**	0.008	0.049**	0.061**	0.007	0.046**	0.067**
	MCP	0.522	0.635**	0.660**	0.535**	0.650**	0.643**	0.523*	0.637**	0.624**
ARIMA (1,1,1)	RMSE	0.016	0.104**	0.109**	0.012	0.067**	0.085**	0.011	0.061**	0.091**
	MAE	0.011	0.078**	0.077**	0.008	0.049**	0.060**	0.007	0.044**	0.063**
	MCP	0.502	0.614**	0.615**	0.482	0.653**	0.659**	0.537**	0.659**	0.652**
ARFIMA(1,d,1)	RMSE	0.017	0.116	0.121	0.013	0.085	0.108	0.012	0.081	0.120
	MAE	0.011	0.087	0.087	0.008	0.063	0.079	0.008	0.059	0.083
	MCP	0.502	0.465	0.501	0.477	0.487	0.501	0.480	0.490	0.498

Table 8 Entries report the values for the root mean squared error (RMSE), mean absolute prediction error (MAE) and the mean correct prediction (MCP) of the direction of change for the different model forecasts over the out-of-sample period from January 4th 2000 to October 30th 2009. The null hypothesis is that the random walk and the model under consideration perform equally well against the alternative that the model under consideration performs better. One and two asterisks denote rejection of the null in favour of the alternative at the 5% and 1% level, respectively.

		30 days			60 days			90 days		
		$MFIV_t$	$SKEW_t$	$KURT_t$	$MFIV_t$	$SKEW_t$	$KURT_t$	$MFIV_t$	$SKEW_t$	$KURT_t$
α = 0.25	**RMSE**	AR, Economic Variables, VECM, ARIMA, Random Walk	VECM, ARIMA	VECM, ARIMA	AR, Economic Variables, VAR, VECM, ARIMA, Random Walk	AR, Economic Variables, VAR, VECM, ARIMA	AR, VAR, VECM, ARIMA	All models	ARIMA	ARIMA
	MAE	AR, Economic Variables, VECM, ARIMA, Random Walk	AR, VAR, VECM, ARIMA	AR, VECM, ARIMA	AR, Economic Variables, VAR, VECM, ARIMA, Random Walk	AR, VAR, VECM, ARIMA	AR, VAR, VECM, ARIMA	AR, Economic Variables, VAR, VECM, ARIMA, Random Walk	ARIMA	ARIMA
α = 0.1	**RMSE**	AR, Economic Variables, VAR, VECM, ARIMA, Random Walk	AR, VECM, ARIMA	AR, VECM, ARIMA	AR, Economic Variables, VAR, VECM, ARIMA, Random Walk	AR, Economic Variables, VAR, VECM, ARIMA	AR, Economic Variables, VAR, VECM, ARIMA	All models	ARIMA	ARIMA
	MAE	AR, Economic Variables, VECM, ARIMA, Random Walk	AR, VAR, VECM, ARIMA	AR, VECM, ARIMA	AR, Economic Variables, VAR, VECM, ARIMA, Random Walk	AR, VAR, VECM, ARIMA	AR, VAR, VECM, ARIMA	AR, Economic Variables, VAR, VECM, ARIMA, Random Walk	ARIMA	ARIMA
α = 0.05	**RMSE**	AR, Economic Variables, VAR, VECM, ARIMA, Random Walk	AR, VAR, VECM, ARIMA	AR, VECM, ARIMA	All models	AR, Economic Variables, VAR, VECM, ARIMA	AR, Economic Variables, VAR, VECM, ARIMA	All models	ARIMA	ARIMA
	MAE	AR, Economic Variables, VECM, ARIMA, Random Walk	AR, VAR, VECM, ARIMA	AR, VECM, ARIMA	AR, Economic Variables, VAR, VECM, ARIMA, Random Walk	AR, VAR, VECM, ARIMA	AR, VAR, VECM, ARIMA	AR, Economic Variables, VAR, VECM, ARIMA, Random Walk	ARIMA	ARIMA

Table 9 Entries report the forecasting models contained in the Model Confidence Set (MCS) related to significance levels α of 0.25, 0.1 and 0.05. For each level of significance, the composition of the MCS for each moment and maturity is reported based on RMSE and MAE, separately. The out-of-sample assessment is conducted over the period from January 4th 2000 to October 30th 2009.

Moneyness	Type of option
Panel A: Moneyness definitions for the skewness trade	
$m \leq 0.95$	OTM put
$m \geq 1.05$	OTM call
Panel B: Moneyness definitions for the kurtosis trade	
$0.95 < m \leq 1.00$	ATM put
$1.00 \leq m < 1.05$	ATM call
$0.90 < m \leq 0.95$	Near OTM put
$1.05 \leq m < 1.1$	Near OTM call
$m \leq 0.9$	Far OTM put
$m \geq 1.1$	Far OTM call

Table 10 Entries report the moneyness m (defined as $m = \frac{K}{S}$) region definitions for the options involved in the risk-neutral skewness (Panel A) and kurtosis trades (Panel B).

Panel A: Trading Profits based on AR(1) Forecasts			
	Sharpe Ratio	**Leland's Alpha**	**S&P 500 Sharpe Ratio**
Trading Skewness	2.45*	0.10*	-0.83
	[1.83, 3.10]	[0.08, 0.13]	[-1.32, -0.34]
Trading Kurtosis	-0.03	-0.13	-0.83
	[-0.61, 0.67]	[-2.72, 2.58]	[-1.32, -0.34]

Panel B: Trading Profits based on Economic Variables Model Forecasts			
	Sharpe Ratio	**Leland's Alpha**	**S&P 500 Sharpe Ratio**
Trading Skewness	2.37*	0.10*	-0.83*
	[1.76, 3.02]	[0.07, 0.12]	[-1.32, -0.34]
Trading Kurtosis	0.01	0.05	-0.83*
	[-0.54, 0.72]	[-2.47, 2.70]	[-1.32, -0.34]

Panel C: Trading Profits based on VAR(1) Forecasts			
	Sharpe Ratio	**Leland's Alpha**	**S&P 500 Sharpe Ratio**
Trading Skewness	2.11*	0.09*	-0.83*
	[1.45, 2.80]	[0.06, 0.12]	[-1.32, -0.34]
Trading Kurtois	0.01	0.05	-0.83*
	[-0.54, 0.71]	[-2.59, 2.61]	[-1.32, -0.34]

Panel D: Trading Profits based on VECM(1) Forecasts			
	Sharpe Ratio	**Leland's Alpha**	**S&P 500 Sharpe Ratio**
Trading Skewness	2.20*	0.09*	-0.83*
	[1.64, 2.77]	[0.06, 0.12]	[-1.32, -0.34]
Trading Kurtosis	0.21	0.86	-0.83*
	[-0.39, 0.93]	[-1.90, 3.57]	[-1.32, -0.34]

Panel E: Trading Profits based on ARIMA(1,1,1) Forecasts			
	Sharpe Ratio	**Leland's Alpha**	**S&P 500 Sharpe Ratio**
Trading Skewness	2.47*	0.10*	-0.83*
	[1.92, 3.08]	[0.08, 0.13]	[-1.32, -0.34]
Trading Kurtosis	-0.26	-1.34	-0.83*
	[-0.69, 0.33]	[-4.18, 1.25]	[-1.32, -0.34]

Panel F: Trading Profits based on Random Walk Forecasts			
	Sharpe Ratio	**Leland's Alpha**	**S&P 500 Sharpe Ratio**
Trading Skewness	-0.64*	-0.03*	-0.83*
	[-1.01, -0.28]	[-0.04, -0.01]	[-1.32, -0.34]
Trading Kurtosis	-0.57*	-3.01*	-0.83*
	[-1.04, -0.004]	[-6.60, -0.10]	[-1.32, -0.34]

Table 11 Entries report the annualised Sharpe Ratio and Leland's Alpha for the skewness and kurtosis trading strategy conducted from January 4th 2000 to October 30th 2009. Different model forecasts have been used to form the trading strategies on (Panels A to F). The bootstrapped 95% confidence intervals for the Sharpe Ratio and Leland's alpha are provided in brackets. The S&P 500 buy & hold Sharpe Ratio is also reported. One asterisk indicates that the reported figure is statistically different from zero.

Panel A: Trading Profits based on AR(1) Forecasts		
	Sharpe Ratio	**Leland's Alpha**
Trading Skewness	-30.44*	-1.78*
	[-33.27, -28.20]	[-1.87, -1.70]
Trading Kurtosis	-2.93*	-36.07*
	[-5.20, -2.36]	[-45.29, -28.71]

Panel B: Trading Profits based on Economic Variables Model Forecasts		
	Sharpe Ratio	**Leland's Alpha**
Trading Skewness	-30.66*	-1.79*
	[-33.42, -28.48]	[-1.87, -1.71]
Trading Kurtosis	-2.91*	-35.80*
	[-5.06, -2.34]	[-44.80, -28.46]

Panel C: Trading Profits based on VAR(1) Forecasts		
	Sharpe Ratio	**Leland's Alpha**
Trading Skewness	--30.43*	-1.79*
	[-33.20, -28.18]	[-1.88, -1.71]
Trading Kurtois	-2.91*	-35.63*
	[-5.12, -2.33]	[-44.14, -28.23]

Panel D: Trading Profits based on VECM(1) Forecasts		
	Sharpe Ratio	**Leland's Alpha**
Trading Skewness	-30.37*	-1.80*
	[-33.11, -28.10]	[-1.89, -1.72]
Trading Kurtosis	-2.94*	-34.69*
	[-6.31, -2.32]	[-43.28, -27.25]

Panel E: Trading Profits based on ARIMA(1,1,1) Forecasts		
	Sharpe Ratio	**Leland's Alpha**
Trading Skewness	-29.75*	-1.80*
	[-32.40, -27.72]	[-1.89, -1.71]
Trading Kurtosis	-3.00	-37.84
	[-4.74, -2.43]	[-47.08, 30.02]

Panel F: Trading Profits based on Random Walk Forecasts		
	Sharpe Ratio	**Leland's Alpha**
Trading Skewness	-31.70*	-1.93*
	[-34.35, -29.56]	[-2.03, -1.86]
Trading Kurtosis	-1.19*	-40.34*
	[-4.09, -1.04]	[-63.28, -24.65]

Table 12 Entries report the annualised Sharpe Ratio and Leland's alpha for the skewness and kurtosis trading strategy implemented from January 4th 2000 to October 30th 2009 by including transaction costs. Different model forecasts have been used to form the trading strategies on (Panels A to F). The bootstrapped 95% confidence intervals for the Sharpe Ratio and Leland's alpha are provided in brackets. The S&P 500 buy & hold Sharpe Ratio is also reported. One asterisk indicates that the reported figure is statistically different from zero.

		30 days			60 days			90 days		
		$MFIV_t$	$SKEW_t$	$KURT_t$	$MFIV_t$	$SKEW_t$	$KURT_t$	$MFIV_t$	$SKEW_t$	$KURT_t$
Random Walk	RMSE	0.022	0.094	0.113	0.018	0.083	0.12	0.020	0.072	0.120
	MAE	0.016	0.073	0.087	0.013	0.061	0.086	0.011	0.047	0.076
	MCP	0.500	0.500	0.500	0.500	0.500	0.500	0.500	0.500	0.500
Economic Variables	RMSE	0.0223	0.087**	0.010**	0.018	0.077**	0.111**	0.017	0.067	0.112*
	MAE	0.016	0.067**	0.076**	0.012	0.057**	0.080**	0.011	0.045*	0.072*
	MCP	0.588**	0.625**	0.628**	0.560**	0.610**	0.602**	0.573**	0.618**	0.597**
AR(1)	RMSE	0.0222**	0.083**	0.095**	0.018	0.076**	0.110**	0.017	0.065*	0.109*
	MAE	0.016	0.064**	0.072**	0.012	0.055**	0.078**	0.011	0.043**	0.071**
	MCP	0.527	0.650**	0.699**	0.544*	0.641**	0.655**	0.559**	0.633**	0.623**
VAR(1)	RMSE	0.0222*	0.083**	0.095**	0.018	0.076**	0.110**	0.017	0.066*	0.110*
	MAE	0.016	0.064**	0.072**	0.012	0.060**	0.077**	0.011	0.044**	0.071*
	MCP	0.527	0.643**	0.682**	0.556*	0.647**	0.641**	0.562**	0.618**	0.596**
VECM(1)	RMSE	0.0222**	0.084**	0.102*	0.018	0.075**	0.109**	0.017	0.065*	0.108*
	MAE	0.015	0.066**	0.079	0.012	0.054**	0.078**	0.011	0.044**	0.072*
	MCP	0.574*	0.625**	0.632**	0.580**	0.645**	0.623**	0.542*	0.628**	0.622**
ARIMA(1,1,1)	RMSE	0.0222	0.091*	0.109*	0.018	0.075**	0.110**	0.017	0.061**	0.101**
	MAE	0.016	0.070**	0.081**	0.013	0.055**	0.079**	0.011	0.042**	0.068**
	MCP	0.521	0.598**	0.610**	0.540*	0.626**	0.632**	0.533	0.639**	0.632**
ARFIMA(1,d,1)	RMSE	0.0228	0.099	0.118	0.019	0.090	0.131	0.017	0.081	0.140
	MAE	0.016	0.077	0.091	0.013	0.067	0.098	0.011	0.054	0.090
	MCP	0.518	0.503	0.587	0.489	0.522	0.513	0.480	0.475	0.470

Table 13 Entries report the values for the root mean squared error (RMSE), mean absolute prediction error (MAE) and the mean correct prediction (MCP) of the direction of change for the different model forecasts over the out-of-sample crisis period from January 3rd 2007 to October 30th 2009. The null hypothesis is that the random walk and the model under consideration perform equally well against the alternative that the model under consideration performs better. One and two asterisks denote rejection of the null in favour of the alternative at the 5% and 1% significance level, respectively.

		30 days			60 days			90 days		
		$MFIV_t$	$SKEW_t$	$KURT_t$	$MFIV_t$	$SKEW_t$	$KURT_t$	$MFIV_t$	$SKEW_t$	$KURT_t$
α = 0.25	**RMSE**	All models	AR, VAR, VECM, ARIMA	AR, VAR, ARIMA	AR, Economic Variables, VECM, ARIMA, Random Walk	VECM, ARIMA	AR, Economic Variables, VAR, VECM, ARIMA	All models	ARIMA	ARIMA
	MAE	Economic Variables, VECM, ARIMA, ARFIMA, Random Walk	AR, VAR, VECM, ARIMA	AR, VAR, ARIMA	AR, Economic Variables, VAR, VECM, ARIMA, Random Walk	AR, VAR, VECM, ARIMA	AR, VAR, VECM, ARIMA	All models	ARIMA	ARIMA
α = 0.1	**RMSE**	All models	AR, VAR, VECM, ARIMA	AR, VAR, VECM, ARIMA, Random Walk	All models	VECM, ARIMA	AR, Economic Variables, VAR, VECM, ARIMA	All models	ARIMA	ARIMA
	MAE	All models	AR, VAR, VECM, ARIMA	AR, VAR, ARIMA	AR, Economic Variables, VAR, VECM, ARIMA, Random Walk	AR, VAR, VECM, ARIMA	AR, VAR, VECM, ARIMA	All models	AR, VAR, VECM, ARIMA	ARIMA
α = 0.05	**RMSE**	All models	AR, Economic Variables, VAR, VECM, ARIMA, Random Walk	AR, VAR, VECM, ARIMA, Random Walk	All models	AR, VAR, VECM, ARIMA	AR, Economic Variables, VAR, VECM, ARIMA	All models	All models	All models
	MAE	All models	AR, Economic Variables, VAR, VECM, ARIMA, Random Walk	AR, VAR, ARIMA	All models	AR, VAR, VECM, ARIMA	AR, VAR, VECM, ARIMA	All models	All models	All models

Table 14 Entries report the forecasting models contained in the Model Confidence Set (MCS) related to significance levels α of 0.25, 0.1, and 0.05. For each level of significance, the composition of the MCS for each moment and maturity is reported based on RMSE and MAE, separately. The out-of-sample assessment is conducted over the period from January 3rd 2007 to October 30th 2009.

Figures

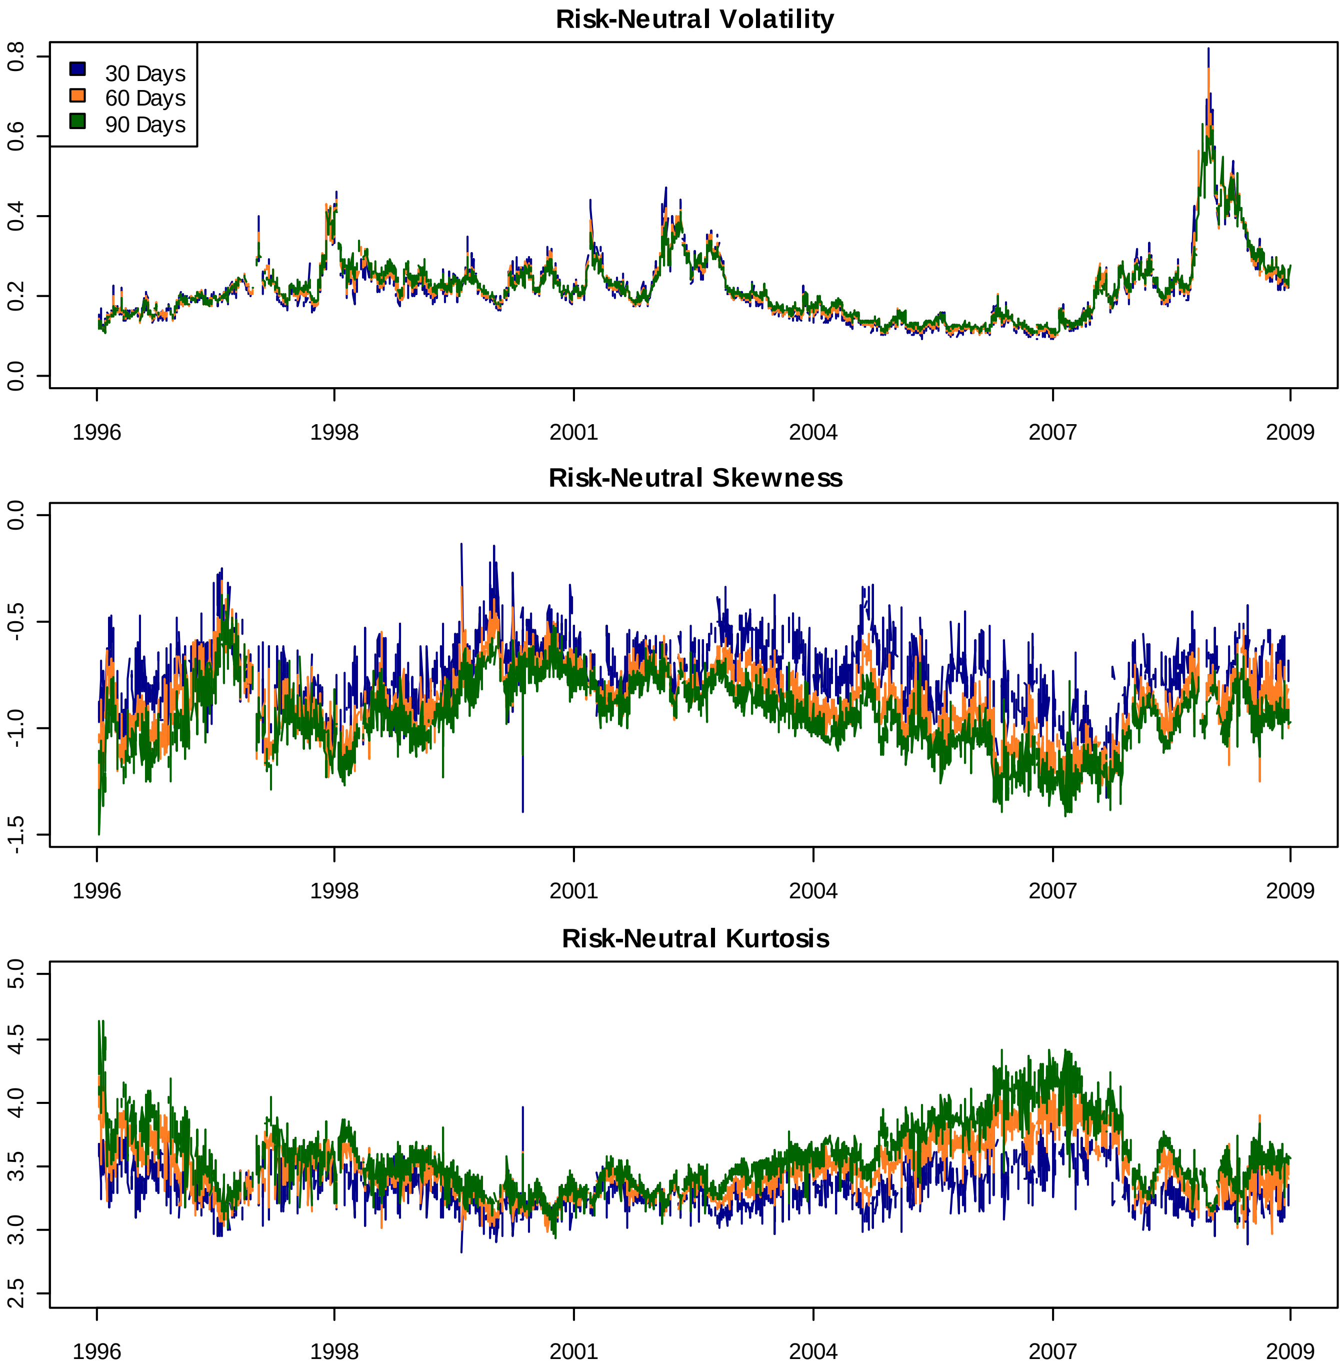

Figure 1 The figure shows the evolution of the 30, 60, and 90-days constant maturity S&P 500 risk-neutral volatility, skewness and kurtosis from January 4th 1996 to October 30th 2009.