JOURNAL OF FINANCIAL AND QUANTITATIVE ANALYSIS Vol. 48, No. 3, June 2013, pp. 947–977
COPYRIGHT 2013, MICHAEL G. FOSTER SCHOOL OF BUSINESS, UNIVERSITY OF WASHINGTON, SEATTLE, WA 98195
doi:10.1017/S002210901300032X

https://doi.org/10.1017/S002210901300032X Published online by Cambridge University Press

Predictable Dynamics in Higher-Order Risk-Neutral Moments: Evidence from the S&P 500 Options

Michael Neumann and George Skiadopoulos*

Abstract

We investigate whether there are predictable patterns in the dynamics of higher-order risk-neutral moments (RNMs) extracted from the market prices of Standard & Poor's (S&P) 500 index options. To this end, we conduct a horse race among alternative forecasting models within an out-of-sample context over various forecasting horizons. We consider both a statistical and an economic setting. We find that higher RNMs can be statistically forecasted. However, only the 1-day-ahead skewness forecasts can be economically exploited. This economic significance vanishes once we incorporate transaction costs. The results have implications for the dynamics of implied volatility surfaces.

I. Introduction

At any given point in time, the cross section of market option prices across strike prices for a given expiry determines the moments of the risk-neutral probability density function (termed risk-neutral moments (RNMs)) of the underlying asset as well as the implied volatility surface (IVS). An IVS maps the implied volatilities extracted from market option prices to different strike prices and maturities. Interestingly, the evolution of RNMs reflects that of IVS and hence that of market option prices. This is because changes in risk-neutral volatility, skewness,

*Neumann, m.neumann@qmul.ac.uk, School of Economics and Finance, Queen Mary, University of London, 327 Mile End Rd, London E1 4NS, UK; Skiadopoulos, gskiado@unipi.gr, Department of Banking and Financial Management, University of Piraeus, Karaoli and Dimitriou 80, Piraeus 18534, Greece, and Queen Mary, University of London, UK, and City University London, UK. We thank Gurdip Bakshi (associate editor and referee) and Paul Malatesta (the editor) for their constructive, stimulating, and thorough comments. We also thank Turan Bali, Alejandro Bernales, Tim Bollerslev, Robert Engle, Daniel Giamouridis, Eric Ghysels, Massimo Guidolin, Erik Hjalmarsson, George Kapetanios, Eirini Konstantinidi, Alexandros Kostakis, Marc Paolella, Paul Schneider, Efthymios Tsionas, Greg Vilkov, and participants at the 2011 Annual Conference of the Society for Financial Econometrics (Chicago), 2011 Annual Conference on Research in Economic Theory and Econometrics (Milos), and 2011 Annual Meeting of the German Finance Association (Regensburg) for useful discussions and comments. Neumann acknowledges financial support from the Society for Financial Econometrics and Skiadopoulos from the Research Centre of the University of Piraeus and the Fundação Ciência e Tecnologia's grant number PTDC/EGE-ECO/099255/2008. Any remaining errors are our responsibility.

and kurtosis are related to the level, slope, and curvature of the IVS, respectively (Zhang and Xiang (2008)).

Understanding the dynamics of the higher-order RNMs is thus of importance to both academics and practitioners. This is because it sheds light on the design of option pricing models that not only price and hedge options accurately at a given point in time but also at *different* points in time. This is a prerequisite for a model to be credible (for the same rationale, see Bakshi, Cao, and Chen (1997) and Dumas, Fleming, and Whaley (1998)). Furthermore, from a trader's perspective, forecasting the changes in RNMs can help in developing profitable risk-neutral volatility, skewness, and kurtosis trading strategies.

We undertake a comprehensive study of the dynamics of higher RNMs extracted from the prices of the Standard & Poor's (S&P) 500 index options, and in particular we investigate whether these can be used to forecast their subsequent changes.[1] To this end, we conduct a horse race among a number of models and evaluate their out-of-sample forecasting performance under a statistical and economic setting. In particular, we apply formal statistical tests to assess whether there are predictable patterns in the dynamics of RNMs, and we determine the set of best-performing models according to the model confidence set (MCS) methodology proposed by Hansen, Lunde, and Nason (2011). Next, we assess the economic significance of the forecasts of the RNMs. We propose and apply risk-neutral skewness and kurtosis option trading strategies based on the forecasted RNMs and evaluate their performance. To address the posed questions, we employ a data set of European index options written on the S&P 500 spanning the period 1996–2010. This is a rich period that includes bull and bear regimes, as well as the recent subprime crisis period; its use is necessary to uncover the dynamics of RNMs given that these depend on relatively rare events, too. To check the robustness of our results, we examine higher RNMs of alternative constant horizons and forecast them over different forecasting horizons (daily, weekly, and monthly). We find that the higher RNMs can be statistically forecasted. However, only the 1-day-ahead skewness forecasts can be economically exploited. This economic significance vanishes once we incorporate transaction costs.

We make three contributions to three respective strands of literature. First, we contribute to the literature on the predictability of IVSs extracted from index options. Its findings are mixed, depending on the evaluation metric and data set. Dumas et al. (1998) find that the dynamics of the S&P 500 IVS are highly unstable under an options hedging setting using weekly data. On the other hand, Gonçalves and Guidolin (2006) find that the S&P 500 IVSs are statistically predictable over daily horizons. These studies fit complex parametric models to IVSs. Instead, studying whether there are predictable patterns in the evolution of RNMs provides an alternative and parsimonious way of investigating whether changes in the IVSs can be forecasted and hence extends the literature on the predictability of IVSs.

[1]This question is distinct from the question of whether RNMs can forecast their realized counterparts as well as other financial and economic variables (see Xing, Zhang, and Zhao (2010), Bakshi, Panayotov, and Skoulakis (2011), Kostakis, Panigirtzoglou, and Skiadopoulos (2011), and Giamouridis and Skiadopoulos (2012)).

https://doi.org/10.1017/S002210901300032X Published online by Cambridge University Press

Furthermore, if RNMs prove to be predictable, this will shed light on the factor(s) contributing to the predictability of IVSs.

Second, we contribute to the literature on the dynamics of higher RNMs extracted from index options. From a theoretical point of view, RNMs are expected to vary over time due to regime shifts in micro and macro fundamentals (see David and Veronesi (2002), Guidolin and Timmermann (2003)). Although there is an extensive empirical literature on the dynamics of the 2nd RNM, that is, the implied volatility (see Harvey and Whaley (1992), Konstantinidi, Skiadopoulos, and Tzagkaraki (2008)), surprisingly little attention has been paid to the evolution of risk-neutral skewness and kurtosis over time per se. To the best of our knowledge, the papers by Panigirtzoglou and Skiadopoulos (2004) and Lynch and Panigirtzoglou (2008) are the closest to ours. Panigirtzoglou and Skiadopoulos identify the factors that drive the dynamics of the S&P 500 implied probability density function and assess their economic significance within a value-at-risk setting. Lynch and Panigirtzoglou provide a detailed description of the evolution of RNMs extracted from FTSE 100 and S&P 500 index options. The authors examine whether cross correlations between moments may be useful to forecast the future moment values. However, the analysis is performed only within sample. Moreover, no other models are considered for comparison purposes. We take a more general approach, as discussed above.

Third, our study complements the vast literature on the efficiency of the S&P 500 option market (see Broadie, Chernov, and Johannes (2009) for a review of the performance of certain index option strategies). Santa-Clara and Saretto (2009) find that the profitability of various index option strategies largely disappears once market frictions are taken into account. We add to this literature by investigating option trading strategies based on higher RNMs with and without transaction costs. The prior literature has mostly studied volatility trading strategies (see Konstantinidi et al. (2008) and references therein). To the best of our knowledge, Bali and Murray (2013) is the only paper that proposes a strategy that allows taking exposure to the time variation of risk-neutral skewness.[2]

The rest of the paper is structured as follows: Section II describes the data set. Section III describes how we extract the RNMs. Section IV describes the forecasting models. Section V presents the in-sample forecasting performance results of these models. Sections VI and VII assess the statistical and economic significance of their out-of-sample performance, respectively, over a daily forecasting horizon. Section VIII investigates the performance of the models over alternative forecasting horizons. Section IX concludes, discusses the implications of the results, and suggests directions for future research.

[2] Ait-Sahalia, Wang, and Yared (2001) and Jha and Kalimipalli (2010) also develop higher-order RNMs strategies, yet they do *not* base them on forecasts for the implied moments per se, as we do in this paper. In the former paper, the strategies are based on deviations between implied moments extracted from the cross section of market option prices and those extracted from the time series of the underlying asset in a risk-neutral world. In the latter paper, the strategies are based on deviations between implied moments and forecasts for the *realized* moments. Interestingly, the Chicago Board Options Exchange (CBOE) introduced in Feb. 2011 an implied skewness index that may serve as the underlying asset to similar type of skewness strategies.

https://doi.org/10.1017/S002210901300032X Published online by Cambridge University Press

II. Data

We obtain daily S&P 500 European style index option data for Jan. 1996–Oct. 2010 from the Ivy DB database of OptionMetrics. For the purposes of our analysis, we use the S&P 500 implied volatilities provided by Ivy DB for each traded contract. We select implied volatilities extracted from options with different strikes and maturities ranging from 5 to 270 days. These are calculated based on the midpoint of bid and ask prices using Merton's (1973) model. In addition, we obtain the closing price of the S&P 500 and the continuously paid dividend yield from Ivy DB.

Following the existing literature on calculating option-implied distributions/RNMs, we impose several filters on the option data set 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 nonzero bid prices and premiums, measured as the midpoint of the best bid and offer, greater than ⅜\$. Third, 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. Finally, we exclude options that form vertical and butterfly spreads with negative prices.

U.S. London Interbank Offered Rates (LIBORs) for maturities of 1–6 months are obtained from Bloomberg as a proxy for the risk-free rate, and the discretely compounded quotes are converted to their continuously compounded counterparts. We obtain rates for any other required maturity by interpolating linearly across the adjacent maturities. In the cases where the desired maturity is beyond the quoted maturities, we extrapolate linearly using the closest available LIBOR.

Finally, we use data for a set of additional economic and financial variables to be used as exogenous predictors in our forecasting models. This set comprises the S&P 500 momentum computed as the previous 6-month S&P 500 return, the implied volatility index VIX downloaded from the Chicago Board Options Exchange (CBOE) Web site (www.cboe.com/micro/VIX/), trading volume of the shortest S&P 500 futures contract taken from Datastream, and trading volume of the S&P 500 options computed from the Ivy DB option data set. We obtain historical data on the 10-year U.S. government bond yields and average yields on Moody's BAA- and AAA-rated corporate bonds from St. Louis Fed Web site (http://research.stlouisfed.org/).

III. Extracting RNMs

We extract RNMs from market option prices using the model-free methodology suggested by Bakshi, Kapadia, and Madan (BKM) (2003).

A. The BKM Method

Let E_t^Q denote the expected value operator under the risk-neutral measure conditional on information 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

https://doi.org/10.1017/S002210901300032X Published online by Cambridge University Press

$R(t,\tau) = \ln[S(t+\tau)] - \ln[S(t)]$ the continuously compounded rate of return at time t over a time period τ. Also, let

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

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

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

denote the fair values of 3 contracts with payoff functions $H[S]$, whereby

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

Furthermore, let

$$\begin{aligned} \mu(t,\tau) &\equiv \mathrm{E}_t^Q\left\{\ln\left[\frac{S(t+\tau)}{S(t)}\right]\right\} \\ &\approx e^{r\tau} - 1 - \frac{e^{r\tau}}{2}V(t,\tau) - \frac{e^{r\tau}}{6}W(t,\tau) - \frac{e^{r\tau}}{24}X(t,\tau) \end{aligned}$$

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 3 contracts, that is,

$$\begin{aligned} \mathrm{MFIV}(t,\tau) &= \sqrt{\mathrm{E}_t^Q\left\{R(t,\tau)^2\right\} - \mu(t,\tau)^2} \\ &= \sqrt{V(t,\tau)e^{r\tau} - \mu(t,\tau)^2}, \end{aligned} \tag{5}$$

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

$$\begin{aligned} \mathrm{KURT}(t,\tau) &= \frac{\mathrm{E}_t^Q\left\{(R(t,\tau) - \mathrm{E}_t^Q[R(t,\tau)])^4\right\}}{\left\{\mathrm{E}_t^Q(R(t,\tau) - \mathrm{E}_t^Q[R(t,\tau)])^2\right\}^2} \\ &= \frac{e^{r\tau}X(t,\tau) - 4\mu(t,\tau)e^{r\tau}W(t,\tau) + 6e^{r\tau}\mu(t,\tau)^2V(t,\tau) - 3\mu(t,\tau)^4}{\left[e^{r\tau}V(t,\tau) - \mu(t,\tau)^2\right]^2}. \end{aligned} \tag{7}$$

The contracts' fair values can be determined by spanning their payoffs by a portfolio of call and put options, as well as the underlying asset and a risk-free bond.

https://doi.org/10.1017/S002210901300032X Published online by Cambridge University Press

It follows that $V(t,\tau)$, $W(t,\tau)$, and $X(t,\tau)$ can be determined by a linear combination of OTM calls and puts, that is,

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

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

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

B. Empirical Implementation

The implementation of equations (8)–(10) requires a continuum of OTM call and put options across strikes. However, market option quotes are available only for a bounded finite range of discrete strike prices. This incurs a bias in the calculation of RNMs (see Dennis and Mayhew (2002), Jiang and Tian (2005)). In addition, we need to extract *constant maturity* moments to eliminate the effect of the shrinking time to maturity on the dynamics of RNMs as time goes by.

To address both issues, once we apply the data filters described in Section II to any given day, we extract the expirations for which at least 2 OTM puts and 2 OTM calls are traded. We discard maturities that do not satisfy this requirement. 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, in line with Malz (1997), we interpolate across the implied volatilities to obtain a continuum of implied volatilities as a function of delta.

https://doi.org/10.1017/S002210901300032X Published online by Cambridge University Press

In particular, 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). We calculate deltas by using the ATM-implied volatility to ensure that the ordering of deltas is the same as the ordering of strike prices. We discard options with deltas above 0.99 and below 0.01, as these correspond to far OTM options that are not actively traded. We also ensure 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 smallest available delta, we extrapolate horizontally using the respective boundary-implied volatility.

To construct the constant maturity moments, we proceed in 4 steps. First, we choose 9 delta values $(0.1, 0.2, \ldots, 0.9)$, and for each we interpolate across the implied volatilities of the various expirations by using a cubic smoothing spline. Then, from the resulting 9 interpolated volatility term structures, we select the respective implied volatilities for a targeted expiration. Next, we obtain the constant maturity-implied volatility curve by fitting a cubic spline through these 9 selected implied volatilities. If the target expiration is below the smallest available traded expiration, a constant maturity-implied volatility curve is not constructed; extrapolating in the time dimension domain yields time series of implied moments that exhibit artificially created spikes. Finally, we convert the delta grid and the corresponding constant maturity-implied volatilities to the associated strike and option prices, respectively, via Merton's (1973) model. Then, we compute the constant maturity moments (eqs. (5)–(7)) by evaluating the integrals in formulae (8)–(10) using trapezoidal approximation. We extract 30-, 60-, and 90-day constant maturities to be used subsequently to verify the robustness of our results.

C. RNMs: A Preliminary Analysis

Figure 1 shows the evolution of the 30-, 60-, and 90-day constant maturity S&P 500 RNMs over the sample period. The risk-neutral volatility is annualized. 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. Both findings are consistent with prior literature documenting that the index risk-neutral probability density functions are negatively skewed and exhibit excess kurtosis (see, e.g., BKM (2003)).

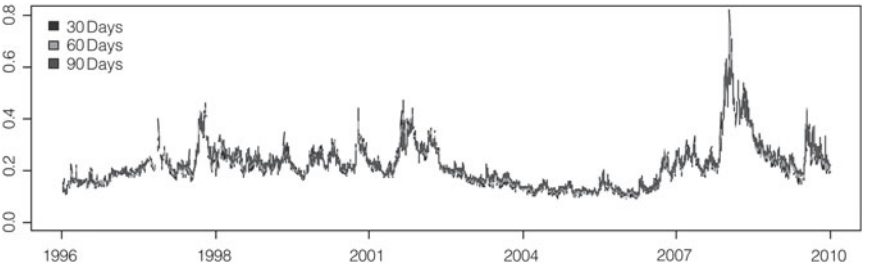
For the purposes of our analysis, we divide the sample into an in-sample part from Jan. 1996–Jan. 2000 and an out-of-sample part spanning the remainder of the data set. Table 1 reports the descriptive statistics of the S&P 500 RNMs measured in levels (Panel A) and 1st differences (Panel B) over the in-sample period. The 1st-order autocorrelation and values of the augmented Dickey-Fuller (ADF) test statistic are reported. We can see that every RNM is positively autocorrelated in the levels, whereas the risk-neutral skewness and kurtosis are negatively autocorrelated when measured in 1st differences. In addition, the unit root test results reveal that all moments (except for the 30-day implied kurtosis) are integrated of order 1. Hence, we will employ the 1st differences throughout our econometric analysis whenever a model requires a stationary time series.

https://doi.org/10.1017/S002210901300032X Published online by Cambridge University Press

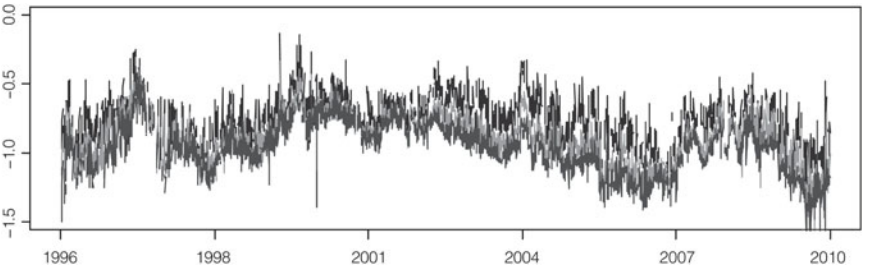
FIGURE 1
Evolution of RNMs over Time

Figure 1 shows the daily evolution of the 30-, 60-, and 90-day constant maturity S&P 500 risk-neutral volatility, skewness, and kurtosis over the period Jan. 4, 1996–Oct. 29, 2010.

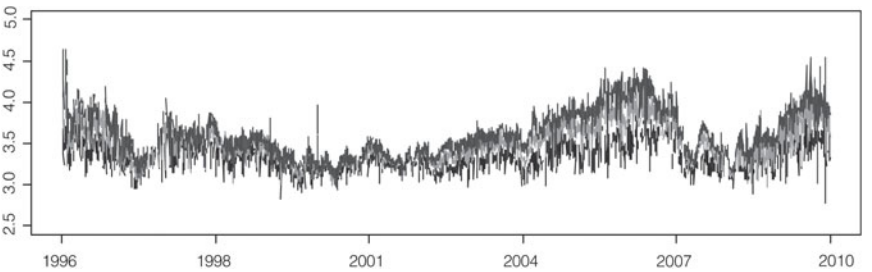
Graph A. Risk-Neutral Volatility



Graph B. Risk-Neutral Skewness



Graph C. Risk-Neutral Kurtosis



D. RNMs and Risk Premiums

Finally, we examine the relationship between the higher RNMs and the variance risk premium (VRP); VRP is defined as the difference between the risk-neutral and statistical expectations of the future return variation (see Bakshi and Madan (2006), Carr and Wu (2009), and Bollerslev, Tauchen, and Zhou (2009) for evidence on the VRP extracted from index options). Exploring this relationship is interesting for two reasons. First, Bakshi and Madan show theoretically that the VRP is related to the physical skewness and kurtosis. Based on this result, Kang, Kim, and Yoon (2010) show theoretically that the VRP is positively related to the risk-neutral skewness and kurtosis. Therefore, confirming that the VRP depends

https://doi.org/10.1017/S002210901300032X Published online by Cambridge University Press

TABLE 1

Descriptive Statistics of S&P 500 RNMs over the In-Sample Period

Table 1 reports descriptive statistics of the extracted S&P 500 risk-neutral moments (RNMs) MFIV, SKEW, and KURT over the in-sample period Jan. 4, 1996–Jan. 3, 2000. Here, ρ_1 denotes the autocorrelation at the 1st lag and ADF the value of the augmented Dickey-Fuller test statistic; * and ** indicate rejection of the null hypotheses of zero 1st-order autocorrelation and existence of a unit root at the 5% and 1% significance levels, respectively. Figures are reported separately for the levels of RNMs (Panel A) and 1st differences (Panel B).

	Dependent Variable								
	30 Days			60 Days			90 Days		
	MFIV	SKEW	KURT	MFIV	SKEW	KURT	MFIV	SKEW	KURT
Panel A. Statistical and Time-Series Properties of RNMs Levels									
No. of obs.	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
Std. 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 RNMs 1st Differences									
No. of obs.	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
Std. 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	−3.50*	−7.50**	−7.50**	−16.80**	−12.00**	−11.80**	−5.20**	−15.30**	−15.00**

on the RNMs empirically will corroborate the results in Bakshi and Madan (2006) and Kang et al. (2010). Second, Kozhan, Neuberger, and Schneider ((2013), Tabs. 6 and 10) find that the VRP proxies the skewness risk premium. The profit from buying skewness swaps is highly correlated with the profit from writing variance swaps.[3] Hence, testing the relationship between the VRP and risk-neutral skewness and kurtosis can shed light on the determinants of the skewness risk premium.

To explore the relationship between RNMs and VRP, first we compute the $\text{VRP}_{t,T}$ at any point t measured over horizons $T = 30$, 60, and 90 days for the period Jan. 2000–Oct. 2010. We do this in order to match the VRP horizon with the corresponding horizon of the extracted risk-neutral skewness $\text{SKEW}_{t,T}$ and kurtosis $\text{KURT}_{t,T}$ (see the Appendix for details on the calculation of $\text{VRP}_{t,T}$). We find that on average, the risk-neutral expectation of the future return variation is greater than the statistical one. This is in line with the findings of the literature on VRP extracted from index options. Subsequently, we run the following regression for each horizon $T \in \{30, 60, 90\}$:

$$\text{VRP}_{t,T} = \beta_0 + \beta_1 \text{SKEW}_{t,T} + \beta_2 \text{KURT}_{t,T} + \varepsilon_{t,T}. \tag{11}$$

[3] A skewness swap is a forward contract that on the expiration date pays to its buyer the difference between the realized (over its life) skewness and the skewness swap rate agreed upon at the inception of the contract. This contract is analogous to a variance swap where the underlying asset is the realized variance instead of the skewness of an asset's returns.

https://doi.org/10.1017/S002210901300032X Published online by Cambridge University Press

Table 2 reports the estimation results for each of the 3 maturities. We can see that in all cases the coefficients of the higher RNMs are positive and statistically significant. This confirms the relationship of the higher RNMs to the VRP and hence to the skewness premium. Furthermore, the results imply that an increase in higher RNMs increases the VRP. This is also consistent with the prediction of the theoretical formula derived by Kang et al. (2010) on the relationship between VRP and higher RNMs.

TABLE 2

VRP and Higher RNMs

Table 2 reports the estimation results of a regression of the variance risk premium ($VRP_{t,T}$) measured over horizon T on the corresponding horizon risk-neutral skewness $SKEW_{t,T}$ and kurtosis $KURT_{t,T}$. The estimated model is

$$VRP_{t,T} = \beta_0 + \beta_1 SKEW_{t,T} + \beta_2 KURT_{t,T} + \varepsilon_t,$$

where $VRP_{t,T}$ denotes the variance risk premium computed as the difference of the expectation of the volatility over T under the risk-neutral and under the physical probability measure. The model is estimated by ordinary least squares (OLS). Separate regressions are performed for horizons $T = \{30, 60, 90\}$ days. The sample spans the period Jan. 3, 2000–Oct. 29, 2010. Newey-West (1987) t-statistics are provided in square brackets, and * and ** indicate rejection of a zero coefficient at the 5% and 1% significance levels, respectively.

	Dependent Variable		
	$VRP_{t,30}$	$VRP_{t,60}$	$VRP_{t,90}$
Incl. obs.	1,760	2,428	2,474
c	−0.94**	−0.83**	−0.61**
	[−6.17]	[−5.88]	[−6.29]
$SKEW_{t,T}$	0.051**	0.36**	0.31**
	[5.79]	[5.27]	[5.07]
$KURT_{t,T}$	0.34**	0.32**	0.24**
	[5.98]	[5.59]	[5.68]
Adj. R^2	0.25	0.30	0.27

IV. The Forecasting Models

We use a number of univariate and multivariate time-series models to examine whether RNMs can be predicted. This is because the question of predictability is a joint hypothesis test of the forecasting model used. The employed models can be grouped into models that only utilize information from the RNM time series per se and models that also exploit exogenous information from economic and financial variables. We consider the latter class of models because various papers document the usefulness of exogenous economic and financial variables for predicting volatility (see Harvey and Whaley (1992), Konstantinidi et al. (2008)) or conditional asymmetry (see Ghysels, Plazzi, and Valkanov (2011)). Next, we motivate the choice of the particular economic and financial variables, and then we describe the models.

A. Economic and Financial Predictor Variables

We employ the S&P 500 index momentum, lagged trading volume in the S&P 500 futures and option market, put/call ratio, VIX, measures of the term spread and default spread, as well as the lagged risk-free rate as exogenous predictors.

https://doi.org/10.1017/S002210901300032X Published online by Cambridge University Press

Regarding the use of the momentum factor as a predictor, Harvey and Siddique (2000) find that when past returns have been high, skewness is forecasted to become more negative. This phenomenon can be understood within a "bubbles theory" setting: High past returns imply that the bubble has been inflating for a long time, and hence there is a larger drop when it bursts and prices revert to fundamentals. The bursting of the bubble will also affect kurtosis because it will yield outliers and create a left fat tail (Blanchard and Watson (1982)). In line with Harvey and Siddique (2000), we use the previous 6 months' S&P 500 return to compute the momentum variable.

We also include the trading volume as a predictor for RNMs. The theoretical model of Hong and Stein (2003) predicts that skewness becomes more negative as the trading volume increases. This is because volume proxies for the intensity of disagreement among market participants, where some of them face short sales restrictions. The empirical evidence in Hansis, Schlag, and Vilkov (2010) corroborates the use of trading volume as predictor of subsequent skewness extracted from individual equity options. In addition, Bollen and Whaley (2004), and Gârleanu, Pedersen, and Poteshman (2009) propose a demand-based option pricing approach and find that measures of the trading volume forecast the shape of the implied volatility curves (i.e., the RNMs). Notice that we include the trading volume of the underlying asset (proxied by the volume of the shortest S&P 500 futures contract) as well as the trading volume of the options as separate predictors. Taylor, Yadav, and Zhang (2009) find that it is only the volume in the equity option markets that helps explain future skewness, whereas the volume of the underlying asset does not. Instead, we prefer taking an agnostic view.

The use of the put/call ratio as a predictor is justified because it is regarded as a measure of the market sentiment that affects RNMs. In particular, in the case where the market is pessimistic, the volume of puts is expected to be greater than the volume of calls (see Cremers and Weinbaum (2010)). This is because investors expect the market to decline and thus the risk-neutral probability density function will appear to be negatively skewed with a pronounced left tail. The use of the VIX is motivated by the volatility feedback theory that implies that an increase in volatility increases the index expected returns and hence decreases the index price. This results in negative skewness and excess kurtosis, too (see Campbell and Hentschel (1992)). The relationship of the market sentiment and volatility measures with risk-neutral skewness and kurtosis is expected to be contemporaneous (see Dennis and Mayhew (2002), Taylor et al. (2009)), yet it may be set within a predictive setting, too. This is because the put/call and VIX variables are persistent, and hence the duration of their effect on higher RNMs is expected not to have a short life.

We further employ the term premium (computed as the difference between the 10-year and 3-month U.S. government bond yields) and the default premium (computed as the difference between the Moody's BAA-rated corporate bonds and the AAA-rated bonds average yields) as predictors for the RNMs. We employ these variables because they serve as leading indicators for the business cycle. Hence, they may prove useful in predicting the RNMs that are also expected to depend on the business cycle (see Harvey and Whaley (1992)).

https://doi.org/10.1017/S002210901300032X Published online by Cambridge University Press

Finally, we also include the lagged risk-free rate as a predictor for the RNMs. This is because RNMs depend on the mean return of the underlying asset, which under the risk-neutral measure is the risk-free rate. Given that the dynamics of interest rates are commonly modeled as an autoregressive process of order 1, the lagged risk-free rate should predict the future RNMs.

B. Univariate ARIMA(X) Models

We employ univariate autoregressive integrated moving average models with (ARIMAX) and without (ARIMA) exogenous regressors to study the dynamics of RNMs. These models capture the autocorrelation in RNMs already documented in Section III.C, as well as autoregressive structures in the error terms commonly attributed to short-term effects. This choice is also consistent with papers that have assumed an autoregressive pattern in the dynamics of the respective physical moments (Harvey and Siddique (1999)). Moreover, the usage of ARIMA models is supported by the existing literature on RNMs. Konstantinidi et al. (2008) find that ARIMA models without exogenous regressors are useful to forecast implied volatilities; we extend their analysis to a higher-order RNMs setting. In addition, Chang, Christoffersen, and Jacobs (2013) document that employing both autoregressive as well as moving average (MA) effects in modeling RNM series is necessary to remove residual autocorrelation.

The ARIMAX(p, d_j, q) model is described by

$$\Phi_j(L)\Delta^{d_j}\text{MOMENT}_{j,t} = c_j + \Theta_j(L)\varepsilon_{j,t} + \sum_{k=1}^{K}\sum_{i=1}^{p}\beta_{k,i}X_{k,t-i}, \tag{12}$$

where d_j denotes the order of differencing required to produce a stationary and invertible process for the jth moment. In our case, $d_j = 1$ for all j, as we find all moments to be integrated of order 1; L denotes the lag operator; $\Phi_j(L) = 1 + \phi_{j,1}L + \cdots + \phi_{j,p}L^p$ and $\Theta_j(L) = 1 + \theta_{j,1}L + \cdots + \theta_{j,p}L^p$ denote the autoregressive and moving average polynomials, respectively; $c_j = -\mu_j(1 + \phi_{j,1} + \cdots + \phi_{j,p})$, with μ_j being the mean of $\Delta\text{MOMENT}_{j,t}$; and K denote the number of exogenous regressors, $X_{k,t}$ the observation of the kth regressor at time t, and $\beta_{k,i}$ the coefficient related to the kth regressor at lag i. The case of setting $\beta_{k,i} = 0$ for all k and i yields an ARIMA$(p, 1, q)$. In our case, the exogenous regressors are the set of variables described in Section IV.A. To avoid overfitting the data, we choose $p = q = 1$ and estimate an ARIMA$(1, 1, 1)$ and an ARIMAX$(1, 1, 1)$ model.

C. Multivariate Models: Vector Error Correction

Next, we consider multivariate time-series models to capture any cross-moment effects in the dynamics of RNMs. Such effects have already been documented in the existing literature on RNMs extracted from index and equity options (see Lynch and Panigirtzoglou (2008), Chang et al. (2013), and Hansis et al. (2010)). From an econometric point of view, one needs to take into account the nonstationarity of the RNMs documented in Section III when formulating a multivariate time-series model for the dynamics of RNMs. This is because the RNMs

https://doi.org/10.1017/S002210901300032X Published online by Cambridge University Press

may be cointegrated; that is, a long-run relationship may exist among them. In the case where there is cointegration, a vector error correction model (VECM) should be used. To investigate whether cointegration exists across RNMs, we apply Johansen's (1988) test for cointegration. We find that the cointegration rank is 2, and we estimate the VECM(1),

$$(13)\quad \begin{pmatrix} \Delta\text{MFIV}_t \\ \Delta\text{SKEW}_t \\ \Delta\text{KURT}_t \end{pmatrix} = A_1 \begin{pmatrix} \Delta\text{MFIV}_{t-1} \\ \Delta\text{SKEW}_{t-1} \\ \Delta\text{KURT}_{t-1} \end{pmatrix} + \Pi \begin{pmatrix} \text{MFIV}_{t-1} \\ \text{SKEW}_{t-1} \\ \text{KURT}_{t-1} \\ 1 \end{pmatrix} + \varepsilon_t,$$

where A_1 is a (3×3) matrix of coefficients controlling for short-run movements in RNMs, and Π is a (3×4) matrix of coefficients controlling the convergence to the long-run equilibrium. Π can be decomposed into a product of (3×2) matrices α and β, and a (1×2) vector c such that $\Pi = \alpha[\beta'\ c']$; $\beta = [\beta_1\ \beta_2]$ contains the two (3×1) cointegrating vectors β_1 and β_2, $\alpha = [\alpha_1\ \alpha_2]$ contains the 2 adjustment vectors, and c contains the constant terms. Given that the decomposition of Π is not uniquely identified, we impose the restrictions $\beta_1 = [1\ 0\ \beta_{13}]'$ and $\beta_2 = [0\ 1\ \beta_{23}]'$. Defining the augmented matrix of cointegrating vectors

$$\beta^* := \begin{bmatrix} \beta \\ c \end{bmatrix},$$

equation (13) becomes

$$(14)\quad \begin{pmatrix} \Delta\text{MFIV}_t \\ \Delta\text{SKEW}_t \\ \Delta\text{KURT}_t \end{pmatrix} = A_1 \begin{pmatrix} \Delta\text{MFIV}_{t-1} \\ \Delta\text{SKEW}_{t-1} \\ \Delta\text{KURT}_{t-1} \end{pmatrix} + \alpha\beta^{*\prime} \begin{pmatrix} \text{MFIV}_{t-1} \\ \text{SKEW}_{t-1} \\ \text{KURT}_{t-1} \\ 1 \end{pmatrix} + \varepsilon_t.$$

Just as in the ARIMAX case, we also employ a variant of the VECM(1) model by adding exogenous predictors. This yields the VECM-X(1) model

$$(15)\quad \begin{pmatrix} \Delta\text{MFIV}_t \\ \Delta\text{SKEW}_t \\ \Delta\text{KURT}_t \end{pmatrix} = A_1 \begin{pmatrix} \Delta\text{MFIV}_{t-1} \\ \Delta\text{SKEW}_{t-1} \\ \Delta\text{KURT}_{t-1} \end{pmatrix} + \alpha\beta^{*\prime} \begin{pmatrix} \text{MFIV}_{t-1} \\ \text{SKEW}_{t-1} \\ \text{KURT}_{t-1} \\ 1 \end{pmatrix} + \Phi X_{t-1} + \varepsilon_t,$$

where X_{t-1} is a (8×1) vector that contains the observations of the economic and financial variables described in Section IV.A measured at time $t-1$, and Φ is the (3×8) coefficient matrix for these variables.

D. Long-Memory ARFIMA Models

Finally, we use a univariate autoregressive fractionally integrated moving average (ARFIMA) model to investigate the presence of long-memory effects. The ARFIMA(p, d_j, q) model is defined as

$$(16)\quad \Phi_j(L)(1-L)^{d_j}(\Delta\text{MOMENT}_{j,t} - \mu_j) = \Theta_j(L)\varepsilon_{j,t},$$

https://doi.org/10.1017/S002210901300032X Published online by Cambridge University Press

where d_j denotes the order of fractional integration of the jth moment, $(1-L)^{d_j}$ the fractional difference operator, and μ_j the expected value of $\Delta\text{MOMENT}_{j,t}$. If $0 < d_j < 0.5$, then the process exhibits "long memory" in the sense that the sum of the autocorrelation functions diverges to infinity. Again, we set $p = q = 1$ to avoid overfitting the data. We perform a maximum likelihood estimation of the ARFIMA$(1, d_j, 1)$ model in the frequency domain by using the Whittle approximation of the Gaussian log-likelihood (see Konstantinidi et al. (2008)). Note that in contrast to the ARIMA$(1, d_j, 1)$, the order of (fractional) integration for the ARFIMA$(1, d_j, 1)$ 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.

Based on the estimated model, we form 1-step-ahead forecasts by taking the infinite autoregressive expansion of the ARFIMA$(1, d_j, 1)$ process, that is,

$$\text{(17)} \qquad \begin{aligned} \text{E}(\text{MOMENT}_{j,t+1}\,|I_t) &= \text{MOMENT}_{j,t} + \mu_j \\ &\quad - \sum_{k=1}^{\infty} \pi_{k,j}(\Delta\text{MOMENT}_{j,t-k+1} - \mu_j), \end{aligned}$$

with

$$\pi_{k,j} = \sum_{i=0}^{k} (b_{i,j} + \phi_j b_{i-1,j})(-\Theta_j)^{k-i} \quad \text{and} \quad b_i = \frac{\Gamma(-d_j + i)}{\Gamma(-d_j)\Gamma(i+1)},$$

where $\Gamma(\cdot)$ denotes the gamma function. To implement equation (17), we truncate the summation at 150.

V. In-Sample Evidence

We first estimate all models described in the previous section using the daily computed RNMs over the in-sample period Jan. 3, 1996–Jan. 3, 2000. Table 3 presents the in-sample estimates for the ARIMA$(1, 1, 1)$ model across all moments and maturities. The t-statistics are reported within square brackets. We can see that the adjusted R^2s are greater for the risk-neutral skewness and kurtosis (ranging from 0.11 to 0.19) compared to these of risk-neutral volatility. Notice that there is evidence of the existence of moving average effects in the RNMs dynamics because the moving average part is statistically significant in most cases. In addition, unreported results of the application of the Durbin-Watson test to the residuals of the estimated ARIMA$(1, 1, 1)$ models reveal that the residuals are uncorrelated for all models and moments. This does not hold uniformly for an AR(1) specification.

Table 4 reports the in-sample estimates for the ARIMAX$(1, 1, 1)$ model. Adding exogenous regressors to the model slightly improves the goodness of fit compared to the ARIMA specification. Regarding the significance of the exogenous regressors, we can see that the futures contract volume and VIX are consistently significant across maturities where risk-neutral volatility is forecasted. When the higher RNMs are forecasted, no predictor is found to be consistently significant across maturities.

https://doi.org/10.1017/S002210901300032X Published online by Cambridge University Press

https://doi.org/10.1017/S002210901300032X Published online by Cambridge University Press

TABLE 3

ARIMA(1, 1, 1): In-Sample Estimation

Table 3 reports the in-sample estimation results of the ARIMA(1,1,1) for each of the 1st differences of the RNMs. The model is

$$(1 + \phi_j L)\,\Delta\text{MOMENT}_{j,t} \quad = \quad c_j + (1 + \theta_j L)\,\varepsilon_{j,t},$$

where $\Delta\text{MOMENT}_{j,t}$ is either $\Delta\text{MFIV}_{j,t}$, $\Delta\text{SKEW}_{j,t}$, or $\Delta\text{KURT}_{j,t}$ and L denotes the lag operator, ϕ_j the autoregressive parameter, and θ_j the moving average parameter. The model is estimated by maximum-likelihood estimation. Results for alternative constant maturity moments with 30-, 60-, and 90-day maturities are provided. The t-statistics are reported in square brackets, and * and ** indicate rejection of a zero coefficient at the 5% and 1% significance levels, respectively. The in-sample period is Jan. 4, 1996–Jan. 3, 2000.

	Dependent Variable								
	30 Days			60 Days			90 Days		
	ΔMFIV_t	ΔSKEW_t	ΔKURT_t	ΔMFIV_t	ΔSKEW_t	ΔKURT_t	ΔMFIV_t	ΔSKEW_t	ΔKURT_t
Incl. obs.	448	448	448	751	751	751	651	651	651
c_j	0.00	0.00	0.00	0.00	0.00*	0.00**	0.00	0.00	0.00
	[1.50]	[0.60]	[−0.17]	[−0.57]	[2.26]	[−2.45]	[−0.21]	[0.43]	[−0.59]
ϕ_j	−0.00	0.22	0.30	−0.51**	−0.15*	−0.13	−0.51**	−0.07	−0.07
	[−0.01]	[1.23]	[1.76]	[−4.38]	[−2.09]	[−1.76]	[−3.16]	[−1.00]	[−0.93]
θ_j	−0.05	−0.13	−0.05	−0.62**	−0.63**	−0.62**	−0.61**	−0.54**	−0.50**
	[−0.07]	[−0.72]	[−0.29]	[−6.01]	[−11.25]	[−10.60]	[−4.11]	[−9.34]	[−8.06]
Adj. R^2	0.00	0.11	0.11	0.01	0.19	0.19	0.01	0.17	0.15

TABLE 4

ARIMAX(1,1,1): In-Sample Estimation

Table 4 reports the in-sample estimation results of the ARIMAX(1, 1, 1) for each of the 1st differences of the RNMs. The model is

$$(1 + \phi_j L)\Delta\text{MOMENT}_{j,t} \quad = \quad c_j + (1 + \theta_j L)\,\varepsilon_{j,t} + \sum_{k=1}^{K} \beta_k X_{k,t-1},$$

where $\Delta\text{MOMENT}_{j,t}$ is either $\Delta\text{MFIV}_{j,t}$, $\Delta\text{SKEW}_{j,t}$, or $\Delta\text{KURT}_{j,t}$ and L denotes the lag operator. The $X_{k,t-1}$ are the exogenous regressors with $k \in \{\text{MOM}_{t-1}, r_{f,t-1}, \text{OPTVOL}_{t-1}, \text{FUTVOL}_{t-1}, \text{PUTCALL}_{t-1}, \text{VIX}_{t-1}, \text{TERM}_{t-1}, \text{DEFAULT}_{t-1}\}$. MOM_{t-1} denotes the level of the S&P 500 momentum, OPTVOL_{t-1} the level of the S&P 500 options trading volume, PUTCALL_{t-1} the S&P 500 put call ratio, VIX_t the level of the CBOE VIX, $r_{f,t}$ the 1st difference of the LIBOR with maturity matching those of the RNMs, FUTVOL_{t-1} the 1st difference of the S&P 500 shortest maturity futures trading volume, TERM_{t-1} the 1st difference of the term premium, and DEFAULT_{t-1} the 1st difference of the default premium. Here, ϕ_j denotes the autoregressive parameter and θ_j the moving average parameter. The model is estimated by maximum-likelihood estimation. Results for alternative constant maturity moments with 30-, 60-, and 90-day maturities are provided. The t-statistics are reported in square brackets, and * and ** indicate rejection of a zero coefficient at the 5% and 1% significance levels, respectively. The in-sample period is Jan. 4, 1996–Jan. 3, 2000.

	Dependent Variable								
	30 Days			60 Days			90 Days		
	ΔMFIV_t	ΔSKEW_t	ΔKURT_t	ΔMFIV_t	ΔSKEW_t	ΔKURT_t	ΔMFIV_t	ΔSKEW_t	ΔKURT_t
Incl. obs.	446	446	446	740	740	740	643	643	643
c_j	−0.01**	0.01	0.00	−0.01**	0.01	−0.01	−0.01**	0.02*	−0.03
	[−3.77]	[0.39]	[−0.05]	[−3.02]	[1.33]	[−1.03]	[−3.05]	[2.02]	[−1.70]
ϕ_j	−0.26	0.16	0.30	−0.47**	−0.14*	−0.13	−0.37	−0.07	−0.07
	[−0.60]	[0.76]	[1.59]	[−3.83]	[−2.04]	[−1.91]	[−1.88]	[−0.96]	[−0.96]
θ_j	−0.38	−0.19	−0.05	−0.59**	−0.65**	−0.65**	−0.47*	−0.54**	−0.51**
	[−0.88]	[−0.86]	[−0.23]	[−5.22]	[−12.42]	[−11.74]	[−2.49]	[−9.62]	[−8.29]
MOM_{t-1}	0.01	0.01	0.00	0.01	−0.02	0.02	0.01**	−0.01	0.00
	[1.36]	[0.16]	[−0.06]	[1.68]	[−0.79]	[0.66]	[2.81]	[−0.60]	[−0.07]
$r_{f,t-1}$	−0.01	0.10	−0.12	0.00	0.08	−0.10	0.00	0.04	−0.05
	[−0.62]	[1.15]	[−1.19]	[0.25]	[1.17]	[−1.10]	[0.07]	[0.36]	[−0.31]
OPTVOL_{t-1}	0.11*	0.17	−0.32	0.02	0.09	−0.17	0.01	−0.11	0.19
	[2.56]	[0.56]	[−0.92]	[0.82]	[0.74]	[−1.08]	[0.25]	[−0.69]	[0.78]

(continued on next page)

TABLE 4 (continued)
ARIMAX(1,1,1): In-Sample Estimation

	Dependent Variable								
	30 Days			60 Days			90 Days		
	ΔMFIV$_t$	ΔSKEW$_t$	ΔKURT$_t$	ΔMFIV$_t$	ΔSKEW$_t$	ΔKURT$_t$	ΔMFIV$_t$	ΔSKEW$_t$	ΔKURT$_t$
FUTVOL$_{t-1}$	−0.02*	−0.10	0.12	−0.01**	−0.01	0.01	−0.01**	0.00	0.00
	[−2.52]	[−1.45]	[1.52]	[−2.64]	[−0.15]	[0.28]	[−2.63]	[0.06]	[0.01]
PUTCALL$_{t-1}$	0.00	−0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	[0.82]	[−1.48]	[0.78]	[0.17]	[−0.99]	[0.38]	[−0.72]	[−0.53]	[−0.04]
VIX$_{t-1}$	0.04**	−0.02	0.04	0.02**	−0.04	0.05	0.02**	−0.07	0.10
	[3.39]	[−0.31]	[0.43]	[2.82]	[−1.53]	[1.54]	[2.92]	[−1.87]	[1.66]
TERM$_{t-1}$	−1.56	5.01	−7.60	−0.23	−2.03	−1.50	−0.31	3.50	−2.59
	[−1.32]	[0.55]	[−0.74]	[−0.34]	[−0.45]	[−0.27]	[−0.48]	[0.66]	[−0.33]
DEFAULT$_{t-1}$	−3.75	2.15	13.99	−0.08	21.52	−21.83	−2.16	46.24	−51.11
	[−0.70]	[0.05]	[0.29]	[−0.02]	[0.95]	[−0.77]	[−0.61]	[1.52]	[−1.14]
Adj. R^2	0.09	0.11	0.11	0.04	0.21	0.21	0.05	0.18	0.16

Table 5 reports the estimation results for the VECM(1) with 2 cointegrating vectors for each maturity. Panel A reports the estimates of the augmented cointegrating vectors β_1^* and β_2^*, and Panel B reports the estimates of the adjustment vectors $\alpha = [\alpha_1\ \alpha_2]$ and the estimated coefficients of matrix A_1. The adjusted R^2 of the VECM(1) specifications for risk-neutral skewness and kurtosis across the various maturities are greater than the ARIMA(X) ones, and range from 0.24 to 0.37. Hence, the VECM captures the time variation in RNMs better than the ARIMA(X) from a goodness of fit standpoint. In addition, most of the elements in the adjustment vectors and the off-diagonal elements of the coefficient matrix A_1 are significant. This implies that both the multivariate structure of the VECM (i.e., cross-moment effects) and cointegration predict RNMs in-sample.

Next, we investigate the incremental effect of including exogenous regressors to the VECM(1) model. Table 6 reports the estimation results for the VECM-X(1). The inclusion of exogenous regressors increases the adjusted R^2 substantially for the risk-neutral volatility for all maturities, whereas the effect on the risk-neutral skewness and kurtosis cases is not uniform across maturities. Similarly to the ARIMAX$(1,1,1)$ model, VIX is the only exogenous predictor that is consistently significant across maturities when it comes to forecasting the risk-neutral volatility.

Finally, Table 7 reports the estimates for the ARFIMA$(1,d_j,1)$ model. We can see that the fractional differencing parameter d_j is significantly different from 0 in only a few cases. Moreover, it lies between 0 and 0.5 for the 60-day risk-neutral kurtosis, indicating that long-memory effects may be present only for this single case. The moving average polynomial parameter θ_j is found to be statistically significant in many cases, just as was the case with the findings for the MA part of the ARIMA$(1,1,1)$ model.

VI. Out-of-Sample Forecasting Performance

We assess the out-of-sample forecasting performance of the models described in Section V over the period Jan. 4, 2000–Oct. 29, 2010. We estimate each model

https://doi.org/10.1017/S002210901300032X Published online by Cambridge University Press

TABLE 5

VECM(1): In-Sample Estimation

Table 5 reports the in-sample estimation results of the VECM(1) for each of the RNMs series. The model is

$$\begin{pmatrix} \Delta\text{MFIV}_t \\ \Delta\text{SKEW}_t \\ \Delta\text{KURT}_t \end{pmatrix} = A_1 \begin{pmatrix} \Delta\text{MFIV}_{t-1} \\ \Delta\text{SKEW}_{t-1} \\ \Delta\text{KURT}_{t-1} \end{pmatrix} + \alpha\beta^{*\prime} \begin{pmatrix} \text{MFIV}_{t-1} \\ \text{SKEW}_{t-1} \\ \text{KURT}_{t-1} \\ 1 \end{pmatrix} + \varepsilon_t.$$

The estimates for the matrix of augmented cointegrating vectors $\beta^* = [\beta_1^* \ \beta_2^*]$ are provided in Panel A. Panel B reports the estimates for the matrix of adjustment vectors $\alpha = [\alpha_1 \ \alpha_2]$ and the coefficient matrix A_1. The cointegrating vectors are estimated by maximum-likelihood estimation, and the adjustment vectors and the coefficient matrix are estimated by OLS. Results for alternative constant maturity moments with 30-, 60-, and 90-day maturities are provided. Newey-West (1987) *t*-statistics are reported in square brackets, and * and ** indicate rejection of a zero coefficient at the 5% and 1% significance levels, respectively. The in-sample period is Jan. 4, 1996–Jan. 3, 2000.

	30 Days		60 Days		90 Days	
Incl. obs.	326		668		553	
Panel A. Estimates for the Cointegrating Vectors						
	β_1^*	β_2^*	β_1^*	β_2^*	β_1^*	β_2^*
MFIV_{t-1}	1	0	1	0	1	0
SKEW_{t-1}	0	1	0	1	0	1
KURT_{t-1}	0.71** [3.79]	0.46** [2.44]	1.80** [6.82]	−1.22** [−3.91]	0.92** [6.93]	−0.59** [−2.83]
c	−2.54** [−4.02]	−0.81 [−1.27]	−6.51** [−7.00]	5.15** [4.71]	−3.53** [−7.42]	3.09** [4.14]
Panel B. Estimates for the Adjustment Vectors and the Coefficient Matrix						
	Dependent Variable					
	ΔMFIV_t	ΔSKEW_t	ΔKURT_t	ΔMFIV_t	ΔSKEW_t	ΔKURT_t	ΔMFIV_t	ΔSKEW_t	ΔKURT_t
α_1'	−0.02 [−1.18]	0.57** [4.31]	−0.92** [−6.46]	0.00 [−0.08]	0.02 [0.23]	−0.37** [−3.31]	−0.01 [−0.92]	0.29** [2.53]	−0.82** [−5.11]
α_2'	−0.03 [−1.44]	0.39** [2.58]	−0.76** [−4.67]	0.00 [−0.09]	−0.04 [−0.52]	−0.25** [−2.52]	−0.01 [−0.76]	0.09 [1.21]	−0.41** [−3.79]
ΔMFIV_{t-1}	−0.05 [−0.83]	−1.54** [−3.37]	1.36** [2.78]	−0.08* [−1.88]	0.78** [2.60]	−1.05** [−2.90]	−0.11** [−2.52]	0.24 [0.64]	−0.56 [−1.06]
ΔSKEW_{t-1}	0.03 [1.31]	−0.72** [−4.47]	0.44** [2.55]	−0.01 [−1.07]	−0.27** [−2.68]	0.00 [0.04]	−0.02* [−1.70]	−0.55** [−4.73]	0.36* [2.21]
ΔKURT_{t-1}	0.02 [0.83]	−0.25* [−1.76]	0.01 [0.04]	−0.02* [−1.71]	0.06 [0.73]	−0.34** [−3.42]	−0.02 [−1.62]	−0.16* [−2.01]	−0.06 [−0.50]
Adj. R^2	0.00	0.31	0.37	0.01	0.24	0.27	0.01	0.24	0.26

recursively by employing a constant rolling window of daily computed RNMs; the 1st-estimation window spans the period Jan. 4, 1996–Jan. 3, 2000. Subsequently, at each point in time, we form 1-day-ahead forecasts of the changes of RNMs. We evaluate the accuracy of the obtained forecasts by 3 commonly used measures: i) the root mean squared prediction error (RMSE), ii) the mean absolute prediction error (MAE), and iii) the mean correct prediction (MCP) of the direction of change. To evaluate the statistical significance of the obtained figures, first, we examine formally whether the forecasts derived from each of the considered models outperform the random walk model that is selected as a benchmark. Next, we identify the *set* of best forecasting models. We implement all tests separately for each of the 3 maturities.

https://doi.org/10.1017/S002210901300032X Published online by Cambridge University Press

https://doi.org/10.1017/S002210901300032X Published online by Cambridge University Press

TABLE 6

VECM-X(1): In-Sample Estimation

Table 6 reports the in-sample estimation results of the VECM(1) for each of the RNMs series. The model is

$$\begin{pmatrix} \Delta\text{MFIV}_t \\ \Delta\text{SKEW}_t \\ \Delta\text{KURT}_t \end{pmatrix} = A_1 \begin{pmatrix} \Delta\text{MFIV}_{t-1} \\ \Delta\text{SKEW}_{t-1} \\ \Delta\text{KURT}_{t-1} \end{pmatrix} + \alpha\beta^{*\prime} \begin{pmatrix} \text{MFIV}_{t-1} \\ \text{SKEW}_{t-1} \\ \text{KURT}_{t-1} \\ 1 \end{pmatrix} + \Phi X_{t-1} + \varepsilon_t.$$

The estimates for the matrix of augmented cointegrating vectors $\beta^* = [\beta_1^* \; \beta_2^*]$ are provided in Panel A. Panel B reports the estimates for the matrix of adjustment vectors $\alpha = [\alpha_1 \; \alpha_2]$ and the coefficient matrices A_1 and Φ. The vector X_{t-1} contains the $t-1$ observations of the exogenous regressors MOM_{t-1}, $r_{f,t-1}$, OPTVOL_{t-1}, FUTVOL_{t-1}, PUTCALL_{t-1}, VIX_{t-1}, TERM_{t-1}, and DEFAULT_{t-1}. MOM_{t-1} denotes the level of the S&P 500 momentum, OPTVOL_{t-1} the level of the S&P 500 options trading volume, PUTCALL_{t-1} the S&P 500 put call ratio, VIX_t the level of the CBOE VIX, $r_{f,t}$ the 1st difference of the LIBOR with maturity matching those of the RNMs, FUTVOL_{t-1} the 1st difference of the S&P 500 shortest maturity futures trading volume, TERM_{t-1} the 1st difference of the term premium, and DEFAULT_{t-1} the 1st difference of the default premium. The cointegrating vectors are estimated by maximum-likelihood estimation, and the adjustment vectors and coefficient matrices are estimated by OLS. Results for alternative constant maturity moments with 30-, 60-, and 90-day maturities are provided. Newey-West (1987) t-statistics are reported in square brackets, and * and ** indicate rejection of a zero coefficient at the 5% and 1% significance levels, respectively. The in-sample period is Jan. 4, 1996–Jan. 3, 2000.

	30 Days		60 Days		90 Days	
Incl. obs.	325		662		546	
Panel A. Estimates for the Cointegrating Vectors						
	β_1^*	β_2^*	β_1^*	β_2^*	β_1^*	β_2^*
MFIV_{t-1}	1.00	0.00	1.00	0.00	1.00	0.00
SKEW_{t-1}	0.00	1.00	0.00	1.00	0.00	1.00
KURT_{t-1}	0.03**	1.19**	0.00	1.02**	−0.05**	1.18**
	[3.13]	[27.87]	[−0.48]	[24.98]	[−4.58]	[18.25]
c	−0.11**	−3.48**	0.02	−2.97**	0.22**	−3.74**
	[−3.02]	[−22.40]	[0.76]	[−18.47]	[4.51]	[−13.89]

Panel B. Estimates for the Adjustment Vectors and Coefficient Matrices

	Dependent Variable								
	ΔMFIV_t	ΔSKEW_t	ΔKURT_t	ΔMFIV_t	ΔSKEW_t	ΔKURT_t	ΔMFIV_t	ΔSKEW_t	ΔKURT_t
α_1'	−0.55**	1.87**	−1.83**	−0.39**	0.74**	−0.67**	−0.09*	0.14	−0.56
	[−11.02]	[4.56]	[−4.15]	[−15.91]	[3.36]	[−2.53]	[−2.32]	[0.35]	[−1.06]
α_2'	0.03*	0.50**	−0.85**	−0.02*	0.22**	−0.48**	0.00	−0.66**	0.41**
	[1.92]	[3.95]	[−6.20]	[−2.08]	[3.19]	[−5.75]	[−0.34]	[−5.87]	[2.66]
ΔMFIV_{t-1}	0.04	−1.91**	1.63**	−0.01	0.72**	−1.11**	0.00	−0.23**	−0.02
	[0.79]	[−4.08]	[3.23]	[−0.21]	[2.34]	[−2.98]	[−0.45]	[−2.83]	[−0.16]
ΔSKEW_{t-1}	−0.02	−0.74**	0.47**	0.00	−0.43**	0.14	0.02**	0.06	−0.15**
	[−0.82]	[−4.70]	[2.77]	[0.13]	[−4.22]	[1.14]	[4.43]	[1.32]	[−2.45]
ΔKURT_{t-1}	−0.02	−0.27*	0.04	−0.01	−0.02	−0.27**	0.00	0.19	−0.26
	[−1.19]	[−1.86]	[0.24]	[−0.78]	[−0.24]	[−2.65]	[0.20]	[1.59]	[−1.58]
MOM_{t-1}	−0.01	0.06	−0.10	0.01**	−0.01	−0.03	0.02**	0.06	−0.15**
	[−1.25]	[0.85]	[−1.21]	[2.57]	[−0.31]	[−0.67]	[4.43]	[1.32]	[−2.45]
$r_{f,t-1}$	0.43	−5.19	13.28	0.01	0.17*	−0.25**	0.00	0.19	−0.26
	[0.23]	[−0.35]	[0.82]	[1.10]	[2.06]	[−2.48]	[0.20]	[1.59]	[−1.58]
OPTVOL_{t-1}	0.07*	0.28	−0.30	0.01	0.29	−0.31	−0.01	−0.29	0.57*
	[1.75]	[0.88]	[−0.89]	[0.50]	[1.62]	[−1.41]	[−0.51]	[−1.43]	[2.04]
FUTVOL_{t-1}	−0.09	−2.01**	1.97**	−0.09*	−0.44	0.46	−0.12**	−0.44	0.58
	[−1.07]	[−2.86]	[2.60]	[−2.26]	[−1.19]	[1.03]	[−2.93]	[−1.04]	[1.01]
PUTCALL_{t-1}	0.00	−0.01	0.01	0.00	0.00	0.00	0.00	0.00	0.00
	[0.76]	[−1.64]	[1.24]	[−0.61]	[−0.78]	[1.06]	[−1.20]	[−1.31]	[1.46]
VIX_{t-1}	0.54**	−1.17**	0.78*	0.37**	−0.48**	0.08	0.26**	−0.21	0.04
	[11.46]	[−3.02]	[1.87]	[17.39]	[−2.54]	[0.36]	[14.12]	[−1.12]	[0.17]
TERM_{t-1}	−1.78	7.44	−6.67	−0.46	−4.09	5.55	−0.62	3.79	−5.25
	[−1.52]	[0.77]	[−0.64]	[−0.78]	[−0.77]	[0.87]	[−1.00]	[0.61]	[−0.61]
DEFAULT_{t-1}	−3.12	26.02	4.32	−1.88	9.35	0.57	0.01	57.21*	−46.44
	[−0.55]	[0.56]	[0.09]	[−0.63]	[0.35]	[0.02]	[0.00]	[1.69]	[−1.00]
Adj. R^2	0.30	0.33	0.37	0.33	0.22	0.25	0.29	0.24	0.27

TABLE 7

ARFIMA(1, d_j, 1): In-Sample Estimation

Table 7 reports the in-sample estimation results of the ARFIMA(1, d_j, 1) for each of the 1st differences of the RNMs. The model is

$$(1 + \phi_j L)(1 - L)^{d_j}(\Delta \text{MOMENT}_{j,t} - \mu_j) = (1 + \theta_j L)\varepsilon_{j,t},$$

where $\Delta \text{MOMENT}_{j,t}$ is either $\Delta \text{MFIV}_{j,t}$, $\Delta \text{SKEW}_{j,t}$, or $\Delta \text{KURT}_{j,t}$; L denotes the lag operator; ϕ_j denotes the autoregressive parameter; θ_j the moving average parameter; and d_j the fractional differencing parameter. The model is estimated by maximum-likelihood estimation. Results for alternative constant maturity moments with 30-, 60-, and 90-day maturities are provided. The t-statistics are reported in square brackets, and * and ** indicate rejection of a zero coefficient at the 5% and 1% significance levels, respectively. The in-sample period is Jan. 4, 1996–Jan. 3, 2000.

	Dependent Variable								
	30 Days			60 Days			90 Days		
	ΔMFIV_t	ΔSKEW_t	ΔKURT_t	ΔMFIV_t	ΔSKEW_t	ΔKURT_t	ΔMFIV_t	ΔSKEW_t	ΔKURT_t
Incl. obs.	448	448	448	751	751	751	651	651	651
d_j	0.37	−0.13*	−0.17**	−0.08*	0.05	0.49**	−0.09*	0.06	0.04
	[1.56]	[−2.07]	[−2.94]	[−2.23]	[0.29]	[4.28]	[−2.24]	[0.50]	[0.39]
ϕ_j	−0.30*	0.45*	0.57**	0.21	−0.15*	0.08	0.36	−0.07	−0.06
	[−2.27]	[2.19]	[3.29]	[0.24]	[−1.95]	[0.89]	[0.68]	[−0.94]	[−0.90]
θ_j	−0.72**	0.23	0.38*	0.22	−0.68**	−0.91**	0.38	−0.59**	−0.54**
	[−5.15]	[0.92]	[1.77]	[0.25]	[−4.39]	[−22.83]	[0.74]	[−5.39]	[−4.77]
Adj. R^2	0.00	0.11	0.12	0.01	0.19	0.19	0.00	0.17	0.15

A. 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. MDM tests the null hypothesis of equal predictive accuracy of a given forecasting model and a benchmark model. Let (e_{1t}, e_{2t}) denote the h-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

$$H_0 : \text{E}[g(e_{1t}) - g(e_{2t})] = 0. \tag{18}$$

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

$$V(\bar{d}) \approx T^{-1}\left[\gamma_0 + 2\sum_{k=1}^{h-1}\gamma_k\right], \tag{19}$$

where γ_k denotes the kth autocovariance of d_t and h the forecasting horizon. Here, γ_k is estimated by

$$\hat{\gamma}_k = T^{-1}\sum_{t=k+1}^{n}(d_t - \bar{d})(d_{t-k} - \bar{d}). \tag{20}$$

https://doi.org/10.1017/S002210901300032X Published online by Cambridge University Press

The null hypothesis (equation (18)) is tested by the MDM test statistic S_1^* given by

$$S_1^* = \left[\frac{T+1-2h+n^{-1}h(h-1)}{T}\right]^{1/2} S_1, \tag{21}$$

where $S_1 = [V(\bar{d})]^{-1/2}\bar{d}$. The critical values of the test are taken from a Student's t-distribution with $(T-1)$ degrees of freedom.

Table 8 reports the values for each of the 3 error metrics for each model, moment, and maturity over the out-of-sample period. Regarding the 1-day-ahead predictability of the risk-neutral volatility, we can see that the VECM-X(1) outperforms the random walk under all metrics. This is in line with the results of Konstantinidi et al. (2008), who find that the random walk is outperformed in forecasting the evolution of VIX. In the case of risk-neutral skewness and kurtosis, all models but the ARFIMA$(1, d_j, 1)$ outperform the random walk in almost all cases.

TABLE 8

Models for Forecasting RNMs: Out-of-Sample Error Metrics

Table 8 reports the values for the root mean squared error (RMSE), mean absolute prediction error (MAE), and mean correct prediction (MCP) of the direction of change for the different model forecasts. The forecasting horizon is 1 day. The models have been estimated on a rolling window with 3 years of daily data. The out-of-sample period is Jan. 4, 2000–Oct. 29, 2010. Results for alternative constant maturity moments with 30-, 60-, and 90-day maturities are provided. The modified Diebold-Mariano test has been applied to test whether each model outperforms the random walk in forecasting a given RNM under a given error metric. 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, and * and ** indicate rejection of the null at the 5% and 1% levels, respectively.

	Dependent Variable								
	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	0.017	0.115	0.127	0.013	0.085	0.110	0.012	0.078	0.117
MAE	0.011	0.085	0.088	0.009	0.060	0.075	0.008	0.053	0.075
MCP	0.500	0.500	0.500	0.500	0.500	0.500	0.500	0.500	0.500
ARIMA(1, 1, 1)									
RMSE	0.017	0.109**	0.119**	0.013	0.075**	0.097**	0.012	0.067**	0.100**
MAE	0.012	0.082**	0.082**	0.008*	0.052**	0.066**	0.008	0.046**	0.066**
MCP	0.525*	0.615**	0.620**	0.541**	0.652**	0.657**	0.528**	0.658**	0.652**
ARIMAX(1, 1, 1)									
RMSE	0.018	0.111**	0.122**	0.014	0.076**	0.099**	0.013	0.068**	0.102**
MAE	0.013	0.081**	0.084**	0.009	0.053**	0.067**	0.008	0.047**	0.067**
MCP	0.479	0.600**	0.628**	0.522*	0.645**	0.640**	0.523*	0.652**	0.652**
VECM(1)									
RMSE	0.017	0.095**	0.101**	0.013	0.074**	0.098**	0.012	0.070**	0.104**
MAE	0.012	0.071**	0.073**	0.008	0.053**	0.067**	0.008	0.048**	0.070**
MCP	0.552**	0.643**	0.670**	0.522*	0.650**	0.639**	0.529**	0.635**	0.626**
VECM-X(1)									
RMSE	0.014**	0.098**	0.107**	0.011**	0.076**	0.102**	0.011**	0.071**	0.108**
MAE	0.012**	0.074**	0.077**	0.007**	0.055**	0.071*	0.007**	0.050**	0.075
MCP	0.709**	0.640**	0.652**	0.695**	0.633**	0.625**	0.649**	0.616**	0.599**
ARFIMA(1, d_j, 1)									
RMSE	0.017	0.109**	0.120**	0.013	0.164	0.120	0.012	0.130	0.100**
MAE	0.012	0.080**	0.082**	0.009	0.062	0.069	0.008	0.049	0.067**
MCP	0.536**	0.618**	0.626**	0.532**	0.651**	0.653**	0.535**	0.654**	0.648**

https://doi.org/10.1017/S002210901300032X Published online by Cambridge University Press

B. Choosing the Best Models: The MCS

In contrast to the MDM test, which tests only whether a given model outperforms a given benchmark, the recently developed MCS methodology by Hansen et al. (2011) selects a set of best forecasting models. Moreover, it acknowledges the information content in the underlying data in the sense that uninformative data yield a larger set 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 $i = 1, \ldots, m_0$. 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 the jth model 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

$$M^* \equiv \{i \in M_0 : \mathrm{E}(d_{ij,t}) \leq 0, \quad \text{for all} \quad j \in M_0\}. \tag{22}$$

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 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:

$$\begin{aligned} H_{0,M} &: \mathrm{E}(d_{ij,t}) = 0, \quad \text{for all} \quad i, j \in M \quad \text{vs.} \\ H_{A,M} &: \mathrm{E}(d_{ij,t}) \neq 0, \quad \text{for some} \quad i, j \in M, \end{aligned} \tag{23}$$

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 MCS.

In order to implement the MCS procedure, one needs an *equivalence test* δ_M to test $H_{0,M}$, as well as an elimination rule e_M that selects the model to be removed from M if necessary. The equivalence test δ_M is based on the range statistic defined as

$$T_M \equiv \max_{i,j \in M} |t_{ij}|, \tag{24}$$

where

$$t_{ij} = \frac{\bar{d}_{ij}}{\sqrt{\mathrm{var}(\bar{d}_{ij})}} \quad \text{with} \quad \bar{d}_{ij} = T^{-1} \sum_{t=1}^{T} d_{ij,t}.$$

If the null hypothesis $H_{0,M}$ is rejected, then the elimination rule given by

$$e_M = \operatorname*{argmax}_{i \in M} \sup_{j \in M} t_{ij} \tag{25}$$

https://doi.org/10.1017/S002210901300032X Published online by Cambridge University Press

selects the worst-performing model in M to be removed from the set of candidate models. The asymptotic distribution of the test statistic T_M under the null hypothesis is nonstandard. Hence, we bootstrap the distribution of the test statistic by implementing the fixed block bootstrap algorithm suggested by Hansen et al. (2011) to estimate the distribution of T_M under the null hypothesis. We set the fixed block length l to a size of 15 observations and construct 2,500 bootstrap resamples. We find that results from the MCS test are robust to alternative choices for l.

We apply the MCS methodology to all RNMs and maturities using both the RMSE and mean squared error (MSE) as the error metrics. Table 9 reports the resulting MCS p-values for all models, RNMs, and maturities. These p-values have to be interpreted in the following way: If the p-value for a given model is greater than some predetermined size of the MCS test α, then this model is included in the α-MCS (see Hansen et al. (2011) for a discussion of the concept of MCS p-values). Two asterisks denote that the model is contained in the 5%-MCS under the respective metric. We can see that even though the composition of the MCS varies across moments and maturities, the random walk is excluded for all investigated maturities and both error metrics. This finding implies that RNMs are predictable and is in line with the findings from the MDM test. Interestingly, in the case where volatility is forecasted, the MCS is a singleton for all maturities containing only the VECM-X(1).

TABLE 9

Models for Forecasting RNMs: MCSs

Table 9 reports the model confidence set (MCS) p-values for the different forecasting models applied to the forecasts for all RNMs under the root mean squared error (RMSE) and mean absolute error (MAE) metric. The forecasting horizon is 1 day. If a p-value for a given model is greater than some predetermined significance level α, then this model is included in the respective α-MCS for the RNM under consideration. Here, ** indicates that the model is contained in the 5%-MCS under the respective metric. Results for alternative constant maturity moments with 30, 60, and 90 days are provided. The models have been estimated on a rolling window with 3 years of daily data. The out-of-sample assessment is conducted over the period Jan. 4, 2000–Oct. 29, 2010.

	Dependent Variable								
	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	0.000	0.000	0.000	0.003	0.000	0.000	0.000	0.000	0.000
MAE	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
ARIMA(1, 1, 1)									
RMSE	0.001	1.000**	1.000**	0.003	0.717**	1.000**	0.000	1.000**	1.000**
MAE	0.000	1.000**	1.000**	0.000	1.000**	1.000**	0.000	1.000**	1.000**
ARIMAX(1, 1, 1)									
RMSE	0.000	0.013	0.001	0.003	0.047	0.045	0.000	0.002	0.003
MAE	0.000	0.491**	0.021	0.000	0.040	0.057**	0.000	0.000	0.020
VECM(1)									
RMSE	0.000	1.000**	1.000**	0.003	1.000**	0.350**	0.000	0.002	0.003
MAE	0.000	1.000**	1.000**	0.000	0.284**	0.057**	0.000	0.000	0.000
VECM-X(1)									
RMSE	1.000**	0.013	0.001	1.000**	0.011	0.024	1.000**	0.002	0.003
MAE	1.000**	0.003	0.000	1.000**	0.002	0.000	1.000**	0.000	0.000
ARFIMA(1, d_j, 1)									
RMSE	0.000	1.000**	0.582**	0.003	0.000	0.024	0.000	0.000	0.590**
MAE	0.000	1.000**	0.377**	0.000	0.002	0.057**	0.000	0.000	0.430**

Our results confirm the findings of Gonçalves and Guidolin (2006), who document that the IVS extracted from the S&P 500 index options is predictable over

https://doi.org/10.1017/S002210901300032X Published online by Cambridge University Press

daily horizons. Moreover, our results show that the documented predictability of the IVS can be attributed not only to the implied volatility but also to the higher RNMs.

VII. Economic Significance

A. Trading Strategies: Design and Implementation

The results on the predictability of the higher RNMs suggest that their dynamics can be forecasted under a statistical setting. To explore the economic significance of the documented predictability in RNMs, we develop option trading strategies (termed risk-neutral or implied skewness and kurtosis strategies) based on the formed out-of-sample risk-neutral skewness and kurtosis forecasts.

We design the proposed strategies based on the relation of the 2 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 (Zhang and Xiang (2008)). 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 BKM (2003)). Therefore, if the risk-neutral skewness is forecasted to decrease, this is equivalent to forecasting that the slope of the implied volatility curve will steepen and vice versa. The strategy will amount to buying OTM puts and selling 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 risk-neutral kurtosis is forecasted to increase, then this is equivalent to forecasting an increase in the curvature of the implied volatility curve. In this case, a strategy would be to sell near-to-the-money options and buy at-the-money and away-from-the-money options (see Ait-Sahalia et al. (2001), Figs. 9 (a) and (b) for an illustration of the skewness and kurtosis trades in the space of risk-neutral probability density functions).

Let SKEW_t and KURT_t denote the risk-neutral skewness and kurtosis, respectively, implied by option prices at time t and $\widehat{\text{SKEW}}_{t,t+1}$ and $\widehat{\text{KURT}}_{t,t+1}$ denote the 1-day-ahead 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: i) If $\text{SKEW}_t > \widehat{\text{SKEW}}_{t,t+1}$, sell all available OTM calls and buy all available OTM puts; ii) if $\text{SKEW}_t < \widehat{\text{SKEW}}_{t,t+1}$, buy all available OTM calls and sell OTM puts; iii) if $\text{SKEW}_t = \widehat{\text{SKEW}}_{t,t+1}$, do nothing. OTM calls (puts) are defined as having strikes that are at least 5% higher (lower) than the underlying price. Similarly, we develop the implied kurtosis trading strategy as follows: i) If $\text{KURT}_t > \widehat{\text{KURT}}_{t,t+1}$, buy all available near to-the-money calls and puts and sell all available ATM and deep OTM calls and puts; ii) if $\text{KURT}_t < \widehat{\text{KURT}}_{t,t+1}$, sell all available near to-the-money calls/puts and buy all available ATM and deep OTM calls and puts; iii) if $\text{KURT}_t = \widehat{\text{KURT}}_{t,t+1}$, do nothing. Deep OTM calls (puts) are defined as having strikes that are more than 10% higher (lower) than the underlying price, and near OTM calls (puts) have strikes that are more than 5% higher (lower) but 10% lower (higher) than the underlying price. ATM calls (puts) have strikes that are higher (lower) than the underlying price but less than 5% away from it.

https://doi.org/10.1017/S002210901300032X Published online by Cambridge University Press

We conduct the skewness and kurtosis trading strategies by employing the 60- and 90-day maturity actual and forecasted implied moments; the 30-day maturity actual and forecasted moments are not used because there are less data for these moments, and hence using them would result in low trading activity. Also, we consider only OTM options because these have greater liquidity than in-the-money options. To immunize the strategies' returns against changes in the underlying asset price and its volatility, we perform both delta and vega hedging (for a similar approach, see Bali and Murray (2013)). 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, we delta-hedge the vega-hedged portfolio by using the S&P 500 as the hedging instrument in line with Bakshi et al. (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 their forecasts) refer to nontraded synthetic constant maturity options. However, for the purposes of implementing the suggested strategies, current and forecasted RNMs that match the maturities of the *traded* options should be employed. To address this mismatch, we select the traded expiration between 60 and 90 days that has the greatest number of option quotes. Next, at any given time t, we interpolate linearly across the 60 and 90 days extracted synthetic RNMs to obtain the (current) RNMs for the selected traded expiration. We also interpolate across the forecasted synthetic RNMs to obtain the forecasted RNMs for the selected traded expiration. Then, we compare these interpolated forecasted RNMs to the interpolated current RNMs and set up the trades. Notice that trading does not take place if there is no available maturity between 60 and 90 days. Similarly, if there is not an extracted RNM at t or its forecast value for $t + 1$, then again, trading does not take place.

B. Performance Evaluation

We evaluate the performance of the trading strategies on the basis of the Sharpe ratio (SR) and Leland's (1999) alpha (A_p). A_p accounts for the presence of (unreported) deviations from normality of the distribution of our strategies returns. It is defined as

$$A_p = \mathrm{E}[R_{\mathrm{STRATEGY}}] - B_p(\mathrm{E}[R_{\mathrm{MARKET}}] - R_{\mathrm{RISKFREE}}) - R_{\mathrm{RISKFREE}}, \tag{26}$$

where R_{STRATEGY}, R_{MARKET}, and R_{RISKFREE} denote the return of the strategy, market portfolio, and the risk-free rate, respectively;

$$B_p = \frac{\mathrm{cov}(R_{\mathrm{STRATEGY}}, -(1 + R_{\mathrm{MARKET}})^{-\gamma})}{\mathrm{cov}(R_{\mathrm{MARKET}}, -(1 + R_{\mathrm{MARKET}})^{-\gamma})}$$

measures systematic risk; and

$$\gamma = \frac{\ln(\mathrm{E}[1 + R_{\mathrm{MARKET}}]) - \ln(1 + R_{\mathrm{RISKFREE}})}{\mathrm{var}[\ln(1 + R_{\mathrm{MARKET}})]}$$

https://doi.org/10.1017/S002210901300032X Published online by Cambridge University Press

measures the relative risk aversion. We approximate the return on the market portfolio by the S&P 500 return, set the risk-free rate equal to the U.S. LIBOR, and estimate A_p in a 2-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:

$$R_{\text{STRATEGY},t} - B_{p,t}(R_{\text{MARKET},t} - R_{\text{RISKFREE},t}) - R_{\text{RISKFREE},t} = A_p + \varepsilon_t. \quad (27)$$

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

Table 10 reports the annualized SR and A_p for the risk-neutral skewness and kurtosis trading strategies over the out-of-sample period Jan. 2000–Oct. 2010. Panels A–E report results for the strategies implemented based on the daily forecasts obtained by the respective models described in Section IV. The SR of the S&P 500 buy-and-hold strategy is also reported. To assess whether the performance measures are statistically different from 0, we provide bootstrapped 95% confidence intervals for the SR and A_p to account for the nonnormality in the strategy returns distribution. We calculate these by employing the Politis and Romano

TABLE 10

Performance Evaluation of the Risk-Neutral Skewness and Kurtosis Strategies: No Transaction Costs

Table 10 reports the annualized Sharpe ratio and Leland's alpha for the skewness and kurtosis trading strategy conducted for Jan. 4, 2000–Oct. 29, 2010. One-day-ahead model forecasts obtained from the respective forecasting models have been used to form the trading strategies (Panels A–E). Trade positions are set up daily and held for 1 day. The bootstrapped 95% confidence intervals for the Sharpe ratio and Leland's alpha are provided in square brackets. The stationary bootstrap with an average block size of 10 has been employed. The S&P 500 buy-and-hold Sharpe ratio is also reported. * indicates that the reported figure is positive and statistically different from zero at the 5% significance level.

	Sharpe Ratio	Leland's Alpha	S&P 500 Sharpe Ratio
Panel A. Performance Based on ARIMA(1, 1, 1) Forecasts			
Trading skewness	2.58* [2.03, 3.15]	0.10* [0.08, 0.12]	−0.18 [−0.60, 0.25]
Trading kurtosis	−0.22 [−0.64, 0.32]	−1.12 [−3.77, 1.36]	−0.18 [−0.60, 0.25]
Panel B. Performance Based on ARIMAX(1, 1, 1) Forecasts			
Trading skewness	2.81* [2.24, 3.39]	0.11* [0.08, 0.13]	−0.18 [−0.60, 0.25]
Trading kurtosis	−0.05 [−0.52, 0.62]	−0.24 [−3.00, 2.38]	−0.18 [−0.60, 0.25]
Panel C. Performance Based on VECM(1) Forecasts			
Trading skewness	1.94* [1.35, 2.50]	0.07* [0.05, 0.10]	−0.18 [−0.60, 0.25]
Trading kurtosis	0.01 [−0.50, 0.69]	0.03 [−2.45, 2.39]	−0.18 [−0.60, 0.25]
Panel D. Performance Based on VECM-X(1) Forecasts			
Trading skewness	2.16* [1.59, 2.77]	0.08* [0.06, 0.11]	−0.18 [−0.60, 0.25]
Trading kurtosis	0.52 [−0.14, 1.31]	2.13 [−0.65, 4.97]	−0.18 [−0.60, 0.25]
Panel E. Performance Based on ARFIMA(1, d_j, 1) Forecasts			
Trading skewness	2.41* [1.87, 2.98]	0.09* [0.07, 0.11]	−0.18 [−0.60, 0.25]
Trading kurtosis	0.09 [−0.52, 0.67]	0.48 [−2.67, 3.60]	−0.18 [−0.60, 0.25]

https://doi.org/10.1017/S002210901300032X Published online by Cambridge University Press

(1994) stationary bootstrap method. We set the average block size to 10 given the (unreported) low autocorrelation of the strategy returns. We also test alternative average block sizes, and results prove to be robust.

We can see that in the case of trading risk-neutral skewness, all models generate economically significant returns because the SRs and A_p are positive and significantly different from 0. This extends the finding reported in Section VI, where these models outperform the random walk under the statistical setting. This is because a trading strategy based on random walk forecasts is equivalent to not trading given the trading rules described in Section VII.A. Moreover, all models outperform the S&P 500 buy-and-hold strategy, which yields an SR of −0.83. Interestingly, the SRs and A_p are quite similar across models. This may be attributed to the fact that the values for the MCP metric in Table 8 are also quite similar across 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 respective moments. With respect to trading risk-neutral 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 risk-neutral skewness yields economically significant profits questions the efficiency of the S&P 500 options market. To fully explore this, we analyze the profitability of the trading strategies proposed in Section VII.A by incorporating transaction costs. In particular, we implement the trading strategies once more using the quoted bid and ask option prices provided by Ivy DB. Table 11 reports the performance metrics for the 2 trading strategies once transaction costs are incorporated. The SR and A_p are negative for both the implied 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 index options market cannot be rejected. This extends the results in Gonçalves and Guidolin (2006), who find that the economic significance of implied volatility trading strategies in the S&P 500 options market over a 1-day horizon vanishes as soon as transaction costs are incorporated.

VIII. Further Robustness Checks: Alternative Forecasting Horizons

In this section, we investigate whether the evidence on the predictability of RNMs prevails over longer forecasting horizons. In particular, we consider weekly and monthly horizons. In order to do so, we form k-step-ahead forecasts by estimating the forecasting models described in Section IV with daily data and iteratively forming k-step-ahead forecasts from these. At time t the forecasting model is estimated using observations from $t - l$ up to t, where l denotes the length of the estimation window. Then, 1-step-ahead forecasts for the RNMs to prevail at $t + 1$ are formed. This forecast value is then plugged into the estimated forecast equation to obtain the $t + 2$ forecast. This procedure is repeated until the forecast for $t + k$ has been obtained. A remark is in order at this point. Despite its simplicity, the iterated forecasts approach poses one challenge with regard to the models that use exogenous predictors. At time t, the information regarding the value of the exogenous predictors is known only up to time t. However, in order

https://doi.org/10.1017/S002210901300032X Published online by Cambridge University Press

TABLE 11

Performance Evaluation of the Risk-Neutral Skewness and Kurtosis Strategies: Transaction Costs

Table 11 reports the annualized Sharpe ratio and Leland's alpha for the skewness and kurtosis trading strategy implemented for Jan. 4, 2000–Oct. 29, 2010 by including transaction costs. One-day-ahead model forecasts obtained from the respective forecasting models have been used to form the trading strategies (Panels A–E). Trade positions are set up daily and held for 1 day. The bootstrapped 95% confidence intervals for the Sharpe ratio and Leland's alpha are provided in square brackets. The stationary bootstrap with an average block size of 10 has been employed. The S&P 500 buy-and-hold Sharpe ratio is also reported.

	Sharpe Ratio	Leland's Alpha
Panel A. Performance Based on ARIMA(1, 1, 1) Forecasts		
Trading skewness	−23.97 [−26.04, −22.13]	− 1.61 [−1.72, −1.50]
Trading kurtosis	− 3.02 [−4.77, −2.45]	−36.24 [−44.71, −29.17]
Panel B. Performance Based on ARIMAX(1, 1, 1) Forecasts		
Trading skewness	−23.74 [−25.90, −21.95]	− 1.60 [−1.72, −1.50]
Trading kurtosis	− 2.88 [−4.22, −2.41]	−38.60 [−47.39, −30.28]
Panel C. Performance Based on VECM(1) Forecasts		
Trading skewness	−24.21 [−26.60, −22.24]	− 1.61 [−1.72, −1.51]
Trading kurtosis	− 2.98 [−6.05, −2.35]	−33.89 [−41.68, −27.36]
Panel D. Performance Based on VECM-X(1) Forecasts		
Trading skewness	−24.22 [−26.63, −22.34]	− 1.62 [−1.73, −1.51]
Trading kurtosis	− 2.97 [−6.50, −2.25]	−31.64 [−39.13, −26.12]
Panel E. Performance Based on ARFIMA(1, d_j, 1) Forecasts		
Trading skewness	−23.87 [−25.93, −22.05]	− 1.62 [−1.72, −1.51]
Trading kurtosis	− 3.07 [−5.08, −2.44]	−35.98 [−44.11, −29.00]

to form the intermediate RNM forecasts up to time $t + k$, the intermediate values for the exogenous predictors are also required. Hence, these need to be forecasted as well. Given that any forecasting specification for the economic/financial predictors would be ad hoc, we prefer not to include the ARIMAX and VECM-X models in the longer forecasting horizon investigation.

Application of the MDM test for the case of weekly $(k = 5)$ and monthly $(k = 21)$ forecasts yields the following unreported results.[4] Regarding the weekly forecasting horizons, we find that risk-neutral skewness and kurtosis can be forecasted across all maturities; there exist models that significantly outperform the random walk, just as was the case with the daily horizon forecasts. This does not hold for the risk-neutral volatility for which evidence of predictability is rather weak because the random walk is rarely outperformed. Regarding the monthly forecasting horizons, the evidence of predictability is mixed. In particular, we find that all 30-day RNMs are unpredictable. On the other hand, there is evidence

[4]Tables A.1–A.4, found on the journal's Web site (journals.cambridge.org/jfqa), contain results for the analysis over the longer horizons.

https://doi.org/10.1017/S002210901300032X Published online by Cambridge University Press

of predictability in risk-neutral skewness and kurtosis for the 60- and 90-day maturities as the ARFIMA$(1, d_j, 1)$ consistently outperforms the random walk in forecasting these moments over the monthly horizon.

Application of the MCS test for weekly and monthly horizons yields the following results. Regarding the weekly horizons, we find that the MCS results support evidence for predictability of risk-neutral skewness and kurtosis. This is because the risk-neutral skewness and kurtosis 5%-MCS do not contain the random walk for any maturity. The exclusion of the random walk from the MCS was also the case in the daily horizons. On the other hand, risk-neutral volatility is not predictable for any maturity; this is in contrast to the daily horizon case. The composition of the 5%-MCS varies across moments and maturities with the ARIMA$(1, 1, 1)$ model being the only model included in all 5%-MCS. Regarding the monthly forecasting horizons, we find that the evidence for predictability is mixed. In particular, there is no predictability for all RNMs for the 30-day maturity because the 5%-MCS contains the random walk. On the other hand, there is evidence for predictability for the 60- and 90-day risk-neutral skewness and kurtosis. The 5%-MCS does not contain the random walk. Notice that the results for the MCS test confirm the MDM for all RNMs across the various maturities.

Overall, the results show that the higher RNMs are consistently statistically predictable for the 60- and 90-day maturities over weekly and monthly forecasting horizons. These findings imply that the S&P 500 IVS is predictable over longer forecasting horizons, too, and thus extend the evidence in Gonçalves and Guidolin (2006). Interestingly, our results imply that this predictability in IVS may be attributed to the predictability of the higher RNMs in the case of the longer forecasting horizons.

In order to assess whether economically significant profits can be generated by trading the RNM forecasts over longer forecasting horizons, we implement the trading strategies described in Section VII using the 1-week- and 1-month-ahead forecasts. In particular, we generate 1-week- (month-) ahead RNM forecasts on each day in the out-of-sample period. These are compared to today's RNMs, and we set up the skewness and kurtosis trades according to the approach described in Section VII. The position set up on a particular day is closed after a week (month) has elapsed. We investigate the performance of these strategies. We find that neither the skewness nor the kurtosis strategy yields economically significant profits over the weekly and monthly horizons even when no-transaction costs are considered. These findings hold regardless of the model under scrutiny.

IX. Conclusions

We investigate whether the dynamics of higher-order risk-neutral moments (RNMs) extracted from the S&P 500 index option prices can be used to forecast their subsequent movements. To the best of our knowledge, our study provides a novel, alternative, and parsimonious method of exploring whether the evolution of the S&P 500 implied volatility surface (IVS) over time can be forecasted.

To address our research question, we perform a horse race among alternative model specifications to investigate their statistical and economic *out-of-sample* forecasting performance. To ensure the robustness of our results, we consider

https://doi.org/10.1017/S002210901300032X Published online by Cambridge University Press

alternative constant maturities and forecasting horizons. We find that higher RNMs can be forecasted statistically; interestingly, this is not the case for risk-neutral volatility in the longer forecasting horizons. However, this predictability is economically significant only for the 1-day-ahead skewness forecasts as long as transaction costs are not considered.

Our findings have four implications. First, the reported predictable patterns in the dynamics of higher-order RNMs imply that the evolution of the S&P 500 IVS is also predictable. This corroborates and extends the results of Gonçalves and Guidolin (2006) to longer forecasting horizons. Second, the results indicate that the predictability in the S&P 500 IVS over longer forecasting horizons is attributed to the predictability of higher-order RNMs rather than implied volatility. Third, autocorrelation in higher-order RNMs as well as cross-moment effects and cointegration between RNMs affect their dynamics over time. These findings should be taken into account for the purposes of developing option pricing models that capture the time variation in IVS. Fourth, we provide evidence for the efficiency of the S&P 500 option market given that no economically significant profits can be attained from the higher RNM strategies once transaction costs are included.

This paper has taken a comprehensive analysis of the dynamics of higher RNMs. A possible extension would be to consider option markets written on other underlying indexes and individual equities. Moreover, it would be worth developing models that explain explicitly why higher-order RNMs may change in a predictable fashion from a theoretical point of view. In the interest of brevity, these topics are best left for future research.

Appendix. Calculating the VRP

The variance risk premium $\mathrm{VRP}_{t,T}$ measured at time t over a period T is defined as

$$\mathrm{VRP}_{t,T} := \mathrm{E}_t^Q[\sigma_{t,t+T}^2] - \mathrm{E}_t^P[\sigma_{t,t+T}^2], \tag{A-1}$$

where $\mathrm{E}_t^Q[\sigma_{t,t+T}^2], \mathrm{E}_t^P[\sigma_{t,t+T}^2]$ denote the time t conditional expectations of the return variance realized over a time period T under the risk-neutral and physical probability measures, respectively. The risk-neutral expectation of the future return variance can be interpreted as the variance swap rate (see Carr and Wu (2009) and references therein). Therefore, we estimate $\mathrm{E}_t^Q[\sigma_{t,t+T}^2]$ using the model-free implied volatility formula developed by Jiang and Tian ((2005), eq. (7)); this formula allows for the presence of jumps in the underlying asset price. We compute the risk-neutral expectation of the future return variance for constant maturities $T = 30$, 60, and 90 days by following the steps for the calculation of RNMs described in Section III.B.

To compute $\mathrm{E}_t^P[\sigma_{t,t+T}^2]$ we follow Bollerslev et al. (2009), that is, we estimate $\mathrm{E}_t^P[\sigma_{t,t+T}^2]$ by the realized variance $\mathrm{RV}_{t-T,t}$ over the past T days. We calculate $\mathrm{RV}_{t-T,t}$ by summing the daily realized variances and the squared overnight returns of the S&P 500 over the period $[t - T, t]$. We obtain the data on daily realized variances from the Web site of the Realized Library published by the Oxford-Man Institute of Quantitative Finance (http://www.oxford-man.ox.ac.uk/). The daily realized variance is the sum of squared 5-minute intraday returns over the respective day.

https://doi.org/10.1017/S002210901300032X Published online by Cambridge University Press

References

Ait-Sahalia, Y.; Y. Wang; and F. Yared. "Do Option Markets Correctly Price the Probabilities of Movement of the Underlying Asset?" *Journal of Econometrics*, 102 (2001), 67–110.

Bakshi, G.; C. Cao; and Z. Chen. "Empirical Performance of Alternative Option Pricing Models." *Journal of Finance*, 52 (1997), 2003–2049.

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

Bakshi, G., and D. Madan. "A Theory of Volatility Spreads." *Management Science*, 52 (2006), 1–12.

Bakshi, G.; G. Panayotov; and G. Skoulakis. "Improving the Predictability of Real Economic Activity and Asset Returns with Forward Variances Inferred from Option Portfolios." *Journal of Financial Economics*, 100 (2011), 475–495.

Bali, T. G., and S. Murray. "Does Risk-Neutral Skewness Predict the Cross-Section of Equity Option Portfolio Returns?" *Journal of Financial and Quantitative Analysis*, forthcoming (2013).

Blanchard, O., and M. Watson. "Bubbles, Rational Expectations, and Financial Markets." In *Crises in the Economic and Financial Structure*, P. Wachter, ed. Lexington, MA: Lexington Books (1982).

Bollen, N. P., and R. E. Whaley. "Does Net Buying Pressure Affect the Shape of Implied Volatility Functions?" *Journal of Finance*, 59 (2004), 711–753.

Bollerslev, T.; G. Tauchen; and H. Zhou. "Expected Stock Returns and Variance Risk Premia." *Review of Financial Studies*, 22 (2009), 4463–4492.

Broadie, M.; M. Chernov; and M. Johannes. "Understanding Index Option Returns." *Review of Financial Studies*, 22 (2009), 4493–4529.

Campbell, J. Y., and L. Hentschel. "No News Is Good News: An Asymmetric Model of Changing Volatility in Stock Returns." *Journal of Financial Economics*, 31 (1992), 281–318.

Carr, P., and L. Wu. "Variance Risk Premiums." *Review of Financial Studies*, 22 (2009), 1311–1341.

Chang, B. Y.; P. Christoffersen; and K. Jacobs. "Market Skewness Risk and the Cross-Section of Stock Returns." *Journal of Financial Economics*, 107 (2013), 46–68.

Cremers, M., and D. Weinbaum. "Deviations from Put-Call Parity and Stock Return Predictability." *Journal of Financial and Quantitative Analysis*, 45 (2010), 335–367.

David, A., and P. Veronesi. "Option Prices with Uncertain Fundamentals." Working Paper, University of Chicago (2002).

Dennis, P., and D. Mayhew. "Risk-Neutral Skewness: Evidence from Stock Options." *Journal of Financial and Quantitative Analysis*, 37 (2002), 471–492.

Dumas, B.; J. Fleming; and R. E. Whaley. "Implied Volatility Functions: Empirical Tests." *Journal of Finance*, 53 (1998), 2016–2059.

Gârleanu, N.; L. H. Pedersen; and A. M. Poteshman. "Demand-Based Option Pricing." *Review of Financial Studies*, 22 (2009), 4259–4299.

Ghysels, E.; A. Plazzi; and R. Valkanov. "Conditional Skewness of Stock Market Returns in Developed and Emerging Markets and Its Economic Fundamentals." Working Paper, University of North Carolina at Chapel Hill (2011).

Giamouridis, D., and G. Skiadopoulos. "The Informational Content of Financial Options for Quantitative Asset Management: A Review." In *Handbook of Quantitative Asset Management*, B. Scherer and K. Winston, eds. Oxford University Press (2012), 243–265.

Gonçalves, S., and M. Guidolin. "Predictable Dynamics in the S&P 500 Index Options Implied Volatility Surface." *Journal of Business*, 79 (2006), 1591–1635.

Guidolin, M., and A. Timmermann. "Option Prices under Bayesian Learning: Implied Volatility Dynamics and Predictive Densities." *Journal of Economic Dynamics and Control*, 27 (2003), 717–769.

Hansen, P. R.; A. Lunde; and J. M. Nason. "The Model Confidence Set." *Econometrica*, 79 (2011), 453–497.

Hansis, A.; C. Schlag; and G. Vilkov. "The Dynamics of Risk-Neutral Implied Moments: Evidence from Individual Options." Working Paper, Goethe University, Frankfurt am Main (2010).

Harvey, C. R., and A. Siddique. "Autoregressive Conditional Skewness." *Journal of Financial and Quantitative Analysis,* 34 (1999), 465–487.

Harvey, C. R., and A. Siddique. "Conditional Skewness in Asset Pricing Tests." *Journal of Finance*, 55 (2000), 1263–1295.

Harvey, C. R., and R. E. Whaley. "Market Volatility Prediction and the Efficiency of the S&P 100 Index Option Market." *Journal of Financial Economics*, 31 (1992), 43–73.

Harvey, D.; S. Leybourne; and P. Newbold. "Testing the Equality of Prediction Mean Squared Errors." *International Journal of Forecasting*, 13 (1997), 281–291.

Hong, H., and J. C. Stein. "Differences of Opinion, Short-Sales Constraints, and Market Crashes." *Review of Financial Studies*, 16 (2003), 487–525.

https://doi.org/10.1017/S002210901300032X Published online by Cambridge University Press

Jha, R., and M. Kalimipalli. "The Economic Significance of Conditional Skewness Forecasts in Index Option Markets." *Journal of Futures Markets*, 30 (2010), 378–406.
Jiang, G. J., and Y. S. Tian. "The Model-Free Implied Volatility and Its Information Content." *Review of Financial Studies*, 18 (2005), 1305–1342.
Johansen, S. "Statistical Analysis of Cointegrating Vectors." *Journal of Economic Dynamics and Control*, 12 (1988), 231–254.
Kang, B. J.; T. S. Kim; and S.-J. Yoon. "Information Content of Volatility Spreads." *Journal of Futures Markets*, 30 (2010), 533–558.
Konstantinidi, E.; G. Skiadopoulos; and E. Tzagkaraki. "Can the Evolution of Implied Volatility Be Forecasted? Evidence from European and US Implied Volatility Indices." *Journal of Banking and Finance*, 32 (2008), 2401–2411.
Kostakis, A.; N. Panigirtzoglou; and G. Skiadopoulos. "Market Timing with Option-Implied Distributions: A Forward-Looking Approach." *Management Science*, 57 (2011), 1231–1249.
Kozhan, R.; A. Neuberger; and P. Schneider. "The Skew Risk Premium in the Equity Index Market." *Review of Financial Studies*, 26 (2013), 2174–2203.
Leland, H. E. "Beyond Mean-Variance: Performance Measurement in a Nonsymmetrical World." *Financial Analysts Journal*, 55 (1999), 27–36.
Lynch, D., and N. Panigirtzoglou. "Summary Statistics of Option-Implied Probability Density Functions and Their Properties." Working Paper, Bank of England (2008).
Malz, A. M. "Estimating the Probability Distribution of the Future Exchange Rate from Option Prices." *Journal of Derivatives*, 5 (1997), 18–36.
Merton, R. "Theory of Rational Option Pricing." *Bell Journal of Economics and Management Science*, 4 (1973), 141–183.
Newey, W. K., and K. D. West. "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica,* 55 (1987), 703–708.
Panigirtzoglou, N., and G. Skiadopoulos. "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 (2004), 1499–1520.
Politis, D. N., and J. P. Romano. "The Stationary Bootstrap." *Journal of the American Statistical Association*, 89 (1994), 1303–1313.
Santa-Clara, P., and A. Saretto. "Option Strategies: Good Deals and Margin Calls." *Journal of Financial Markets*, 12 (2009), 391–417.
Taylor, S. J.; P. K. Yadav; and Y. Zhang. "Cross-Sectional Analysis of Risk-Neutral Skewness." *Journal of Derivatives*, 16 (2009), 38–52.
Xing, Y.; X. Zhang; and R. Zhao. "What Does the Individual Option Volatility Smirk Tell Us About Future Equity Returns?" *Journal of Financial and Quantitative Analysis*, 45 (2010), 641–662.
Zhang, J. E., and Y. Xiang. "The Implied Volatility Smirk." *Quantitative Finance*, 8 (2008), 263–284.

https://doi.org/10.1017/S002210901300032X Published online by Cambridge University Press