

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection Lee Kong Chian School Of
Business

Lee Kong Chian School of Business

1-2019

Intraday information from S&P 500 index futures options

Kian Guan LIM

Singapore Management University, kglim@smu.edu.sg

Chen YING

Kian Leong Nelson YAP

Singapore Management University, nelson.yap.2013@pbs.smu.edu.sg

Follow this and additional works at: https://ink.library.smu.edu.sg/lkcsb_research



Part of the [Finance Commons](#), and the [Finance and Financial Management Commons](#)

Citation

LIM, Kian Guan; YING, Chen; and YAP, Kian Leong Nelson. Intraday information from S&P 500 index futures options. (2019). *Journal of Financial Markets*. 42, 29-55.

Available at: https://ink.library.smu.edu.sg/lkcsb_research/6403

This Journal Article is brought to you for free and open access by the Lee Kong Chian School of Business at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection Lee Kong Chian School Of Business by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylds@smu.edu.sg.

Intraday Information from S&P 500 Index Futures Options

Kian Guan Lim ^{*}, Ying Chen [†], and Nelson K.L. Yap [‡]

Revised February 2018

Abstract

In this paper we employ intraday transaction prices of liquid E-mini S&P 500 index futures options to form 10-minutes ahead risk-neutral skewness forecasts and show profitable options trading strategy net of transaction costs. We do not find profitable trading based on 10-minutes ahead risk-neutral volatility and only very marginal cases of profitable trading using kurtosis forecasts. The skewness profitability anomaly may be an indication of informational market inefficiency in intraday S&P 500 futures options markets, which is contrary to findings using longer-span daily and weekly moments. Our results lend credence to the persistence of intraday trading activities in the markets.

JEL classifications

G13; G17; G23

Keywords

Intraday Options trading; Market Efficiency.

^{*}Professor of Quantitative Finance, Singapore Management University

[†]Assoc Professor, Department of Applied Probability and Statistics, National University of Singapore

[‡]PhD candidate, Lee Kong Chian School of Business, Singapore Management University

We acknowledge helpful comments on earlier versions of the paper from faculty participants at the Singapore Management University Symposium, the National University of Singapore RMI Seminar, the International Institute of Forecaster's Korean Conference, the Statistical Computing Asia Conference at Taipei Academia Sinica, and the Australasian Finance and Banking Conference. The insightful comments of the referee are particularly important in a sharper write-up and revision of this paper. We are also grateful to the programming assistance of Chua Wee Song. Funding from the OUB Chair is gratefully acknowledged. Corresponding author: Kian-Guan Lim, Email: kgl@smu.edu.sg, Address: Lee Kong Chian School of Business, 50 Stamford Road, #05-01, SINGAPORE 178899.

Intraday Information from S&P 500 Index Futures Options

Abstract

In this paper we employ intraday transaction prices of liquid E-mini S&P 500 index futures options to form 10-minutes ahead risk-neutral skewness forecasts and show profitable options trading strategy net of transaction costs. We do not find profitable trading based on 10-minutes ahead risk-neutral volatility and only very marginal cases of profitable trading using kurtosis forecasts. The skewness profitability anomaly may be an indication of informational market inefficiency in intraday S&P 500 futures options markets, which is contrary to findings using longer-span daily and weekly moments. Our results lend credence to the persistence of intraday trading activities in the markets.

1 Introduction

We study the intraday dynamics of risk-neutral moments of S&P 500 index futures prices, and test intraday informational market efficiency in the index futures options market. There have been many tests of daily, weekly, and monthly options prices and the general deduction of market efficiency. However, there are relatively few studies on the efficiency of intraday information derived from traded prices of index or index futures options. In this paper we employ information of risk-neutral moments on S&P 500 index futures returns extracted from liquid E-mini index futures options to form 10-minutes ahead forecasts and develop profitable option trading strategies. We find that strategies capturing risk-neutral skewness information are profitable net of transaction costs. We do not find profitable trading based on 10-minutes ahead risk-neutral volatility and only very marginal cases of profitable trading using kurtosis forecasts.

Neumann and Skiadopoulos (2013) use daily S&P 500 index options over January 1996 to October 2010 to extract risk-neutral moments for forecasting and for testing trading strategies over 1-day, 1-week, and 1-month horizons. They find that all the risk-neutral moments can generally be predicted better out-of-sample relative to the random walk benchmark. Using one-day ahead forecasts of the risk-neutral moments to pre-determine option trades, they find that except for one-day ahead skewness forecast, the other moment forecasts do not support profitable trading. After considering transaction cost in the form of the bid-ask spread, they report that skewness forecasts also did not deliver positive

profitability. The results are similar for forecasts involving longer horizons of a week or longer. They thus conclude, “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 one-day horizon vanishes as soon as transaction costs are incorporated.”

On the other hand, Amaya, Christoffersen, Jacobs, and Vasquez (2015) using intraday data from 1993 to 2013 to compute weekly realized variance, skewness, and kurtosis for equity returns, find a very strong negative relationship between realized skewness and the subsequent week’s stock returns. They state that, “A trading strategy that buys stocks in the lowest realized skewness decile and sells stocks in the highest realized skewness decile generates an average weekly return of 19 basis points with a t-statistic of 3.70.” This may or may not be a small positive weekly return rate of a few basis points after transaction costs. Unlike other studies, they did not find any robust and significant relationship between ex ante volatility or ex-ante kurtosis and equity returns. Amaya et.al. are careful to point out that their results could reflect asset pricing model as the implications of ex-ante skewness are on the subsequent cross-section of stock returns. In the asset pricing equilibrium, the results could be interpreted as one of skewness premium when investors commonly hold in aggregate such long-short skewness portfolios and are compensated for risk-bearing. The paper suggested a connection to market efficiency if the time series of the ex-ante skewness could be exploited to produce consistent profitability in trading each of the stocks.

Indeed most published papers to-date concentrated on verifying if ex-ante systematic moments (proxied by aggregation of stocks with different moments into different portfolios) explain cross-sectional differences in ex-post portfolio returns. Conrad, Dittmar, and Ghysels (2013) use daily individual option prices from 1996 to 2005 to infer their underlying stock return risk-neutral moments over horizons of 1-month to 1-year. By forming portfolios ranked by the risk-neutral moments, they find that skewness has a strong negative relation with subsequent returns. Firms with less negative or positive skewness systematically earn lower returns. Bali and Murray (2013) use monthly options data from 1996 to 2010 to construct monthly risk-neutral skewness of stocks to form forward 1-month portfolio of skewness assets abstracted from delta and vega risks, and find a strong negative relation between risk-neutral skewness and the skewness asset returns. As in Conrad, Dittmar, and Ghysels (2013), this is consistent with a positive skewness preference or negative skewness premium in asset pricing theory. Chang, Christoffersen, and Jacobs (2013) estimate market return moments from daily S&P 500 index option data.

They find that risk-neutral market-wide skewness and kurtosis are important risk factors in explaining the cross section of stock returns. Bali, Hu, and Murray (2013), Cremers, Halling, and Weinbaum (2013), Dennis and Mayhew (2002), and Friesen, Zhang, and Zorn (2012) attempt to explain the existence and significance of risk premia related to systematic risk-neutral moments. Thus there are sufficient studies to indicate the importance of examining risk-neutral moments and their relationships to stock expected returns.

However, the role of ex-ante risk-neutral moments in an intertemporal asset pricing model would imply two things; firstly, moments feature as factors in equilibrium expected returns, and secondly, future moments can be forecasted which have risk-adjusted impact on returns. It is the latter aspect of forecasting which we are interested. Moreover, Neumann and Skiadopoulos open up a new line of enquiry into efficiency on the options market itself, not the stock market. Consistency of profitable returns on trading options would be an indication of market inefficiency in the options market.

The risk-neutral probability distribution of the underlying asset to an option embodies a large amount of information on market expectations as well as its risk preferences. In a 2013 NYU Stern–Federal Reserve Conference on Risk Neutral Probability Density, Figlewski (2013) suggested that searching for profitable trading strategies is a good question for research. This is indeed a clever insight as trading profitability not only has obvious attractions for the finance industry, but it also has deep implications on theory. Option theory so far has relied on the risk-neutral probability distribution of the underlying for pricing, so any portfolios of options formed from other options and the underlying should yield a risk-neutral return over a short interval such as 10-minutes. If we can find profitable trading strategies based on pre-determined information, then the options market is informationally inefficient or options theory needs some adjustment.

Some differences in empirical results on profitability after transaction costs could be due to the longer horizon of at least a day if not a week in time series forecasting before deploying the associated options trading strategies. This regards the possibility of information dissipating across time. Another reason for differences could be due to the different methods of forecasting over at least a one-day interval. Yet another difference could be more profitable trading strategies versus strategies that are common knowledge so that they cannot exploit any informational advantages. Amongst practitioners, it is widely known that trading strategies relying on microscopic forecasts of moments or of other variables such as news would best be executed on an intraday basis and not after a whole day or longer. It is thus critical to complement existing studies with a focus on what goes on in intraday dynamics and options trading – what this paper is about. In particular, the option strategies we invoke depend explicitly on the forecasts of future

moments, and our trading profits take into account not just bid-ask costs, but also exogenous transaction costs in the form of additional trading commission costs which were not mentioned in many studies. We avoid micro-structural issues by market orders buying at traded ask prices and selling at traded bid prices, thus absorbing the full impact of the spread cost. We assume our trading volume is small and does not incur impact cost nor compete with high frequency traders.

In our paper, we make a number of contributions to the literature in the area of intra-day futures options trading profitability and intra-day equity futures index return moments predictability. As far as we know, our paper is one of the first studies on intraday implied moments of S&P 500 index futures returns using intraday or high frequency futures option prices. We use intra-day E-mini S&P500 European-style futures options data and improve on existing techniques to extract the first four moments of the risk-neutral return distribution. Secondly we perform intraday out-of-sample forecasting or prediction, and also document the intraday dynamics of the index futures return risk-neutral moments. We introduce a novel local autoregression method that allows variable window in fitting the autoregressive parameters. This is particularly useful in situations when there may be intraday news that cause structural changes in the returns or price distributions. It also distinguishes itself from the conventional autoregressive model with predetermined sample lengths. Thirdly, we show profitability in the trading strategies involving the various risk-neutral moment forecasts, particularly that involving skewness. The positive profitability after transaction costs in skewness trading indicates an anomaly which may be an indication of market inefficiency in intraday markets – it could be due to information inefficiency within short spans of time.

In section 2, we discuss the method for extracting the risk-neutral moments. The data and implementation procedures are then explained. Section 3 provides a discussion of the forecasting models used in the forecast of intraday 10-minutes ahead risk-neutral moments. The local autoregressive model is also explained. Section 4 contains the empirical results showing the forecasting performances of the various models. Section 5 provides the results based on different option trading strategies involving the risk-neutral moment forecasts. Section 6 reports results trying to forecast intraday 10-minutes ahead futures returns based on information of ex-ante risk-neutral moments. Section 7 contains the conclusions.

2 Implied Risk-Neutral Moments

Prediction of returns moments for the purposes of financial trading, hedging, and asset pricing, is prevalent in finance. The most common forecast is that of predictive mean

usually obtained from a regression model. More general frameworks for setting up predictability ranges from the modeling of stochastic models to time series modeling such as GARCH. For forecasting realized volatility, Andersen, Bollerslev, Diebold, and Labys (2003) is a fundamental paper on the employment of high-frequency intraday data. Ghysels, Harvey, and Renault (1996) and Barndorff-Nielsen, Nicolato, and Shephard (2002) provide reviews on stochastic volatility models. Poon and Granger (2003) provide a comprehensive review of forecasting volatility in financial markets including detailed discussion of the pioneering dynamic time series models of Engle (1982) and Bollerslev (1986). Harvey and Siddique (1999, 2000) were some pioneering studies of skewness in financial markets. Brooks, Burke, Heravi, and Persaud (2005) study autoregressive conditional kurtosis. There is a vast literature on empirical measures of asset returns moments, particularly on volatility and skewness, but a relative scarcity of studies on risk-neutral measures of similar moments.

The seminal paper by Breeden and Litzenberger (1978) connected option prices with no-arbitrage state prices, the equivalence of discounted risk-neutral densities. Later papers of the same genre include Rubinstein (1996) and Jackwerth and Rubinstein (1996). Practical approaches of extracting or implying the risk-neutral moments using option prices were developed in Jiang and Tian (2005) and in Bakshi, Kapadia, and Madan (2003). Due to the more general framework and higher moments inferable via the Bakshi, Kapadia, and Madan (2003) method (hereafter referred to as BKM), there has been a surge of interest in using BKM method to study risk-neutral moments. Unless otherwise stated, the research discussed below employ risk-neutral higher moments that are computed using the BKM method.

As for the option prices from which to extract risk-neutral moments, Bakshi, Kapadia, and Madan (2003) and Taylor, Yadav, and Zhang (2009) find that risk-neutral skewness implied from individual stock options are less negative than that implied from stock index option. Gârleanu, Pedersen, and Poteshman (2009) find that recent option-pricing puzzles may be explained by the fact that there is a significant difference between index option prices and the prices of single-stock options due to differences in end-user demands. Figlewski (2008) suggests that the estimation of individual stock risk-neutral density is especially hampered by two serious problems, as stock options trade a relatively small number of strikes, and also face significant microstructural noise. Due to the above, we consider the study of S&P 500 index futures options to be appropriate for the purpose of examining predictability and trading profitability. The S&P 500 index futures options and index futures, as well as index options, are not only more liquid, but also possess more negative skewness moments on the underlying whereby skewness trading strategy

could be effected. For our study we use the S&P 500 index futures and its corresponding futures options.

2.1 BKM Method

The BKM method can be utilized for the case of S&P 500 options on futures. Define the τ -period random log return (or rate of change) of the underlying futures contract at time t as $R(t, \tau) \equiv \ln \left(\frac{F_{t+\tau}}{F_t} \right)$, where F_t is the S&P 500 E-mini index futures price at t . To obtain the risk-neutral variance, skewness and kurtosis of $R(t, \tau)$, it is sufficient to obtain the first 4 risk-neutral moments of $\mathbb{E}^Q[R(t, \tau)]$, $\mathbb{E}^Q[R(t, \tau)^2]$, $\mathbb{E}^Q[R(t, \tau)^3]$, $\mathbb{E}^Q[R(t, \tau)^4]$ under risk-neutral probability measure Q .

Each of the moments above can be viewed as an expected payoff at maturity $t + \tau$ and is a function of some underlying portfolio. Here we rely on a well-known result in Carr and Madan (2001) that any payoff as a function of underlying futures $F(t)$ can be spanned and priced using a traded set of options across different strike prices. For example, a forward can be decomposed as a long call and a short put with same strike. A call spread can be decomposed as a long call with lower strike and a short call with higher strike. For a more complicated payoff, we need more options with different strikes to replicate the payoff. This can be done assuming that the payoff function is twice continuously differentiable in the underlying futures price. We can write $H(j)[F_T] = R(t, \tau)^j$, for $j = 2, 3, 4$. Then consider the expansion of $H(j)[F_T]$:

$$\begin{aligned} H(j)[F_T] &= H(j)[F_t] + (F_T - F_t)H_F(j)[F_t] \\ &\quad + \int_{F_t}^{\infty} H_{FF}[K](F_T - K)^+ dK + \int_0^{F_t} H_{FF}[K](K - F_T)^+ dK \end{aligned}$$

where $H_F(j)$ refers to the derivative of function $H(j)$ with respect to F_T .

Then, letting $V_t(\tau) \equiv E_t^Q(e^{-r\tau}H(2)[F_T])$, where Q is the risk-neutral measure and r is the risk-free rate, we have

$$\begin{aligned} V_t(\tau) &= \int_{F_t}^{\infty} \frac{2(1 - \ln(K/F_t))}{K^2} C_t(\tau; K) dK \\ &\quad + \int_0^{F_t} \frac{2(1 + \ln(F_t/K))}{K^2} P_t(\tau; K) dK \end{aligned}$$

where $C_t(\tau; K)$ and $P_t(\tau; K)$ are respectively the European call and put on the underlying stock index futures price F_t , with maturity τ and strike price K .

The no-arbitrage prices of payoffs at T , $H(3)[F_T]$ and $H(4)[F_T]$ can be written as $W_t(\tau)$ and $X_t(\tau)$, and are similarly found as

$$\begin{aligned}
W_t(\tau) &= \int_{F_t}^{\infty} \frac{6 \ln(K/F_t) - 3(\ln(K/F_t))^2}{K^2} C_t(\tau; K) dK \\
&\quad - \int_0^{F_t} \frac{6 \ln(F_t/K) + 3(\ln(F_t/K))^2}{K^2} P_t(\tau; K) dK \\
X_t(\tau) &= \int_{F_t}^{\infty} \frac{12(\ln(K/F_t))^2 - 4(\ln(K/F_t))^3}{K^2} C_t(\tau; K) dK \\
&\quad + \int_0^{F_t} \frac{12(\ln(F_t/K))^2 + 4(\ln(F_t/K))^3}{K^2} P_t(\tau; K) dK.
\end{aligned}$$

These contract prices $V_t(\tau)$, $W_t(\tau)$, and $X_t(\tau)$ are also the respective discounted risk-neutral moments of the underlying futures return or rate of change over period interval τ . The equations involve weighted sums of out-of-the-money options across varying strike prices. Using these prices, the risk-neutral (central) moments RNMs of variance, skewness, and kurtosis can be calculated as

$$\begin{aligned}
VAR_t^Q(\tau) &= e^{r\tau} V_t(\tau) - \mu_t(\tau)^2 \\
SKEW_t^Q(\tau) &= \frac{e^{r\tau} W_t(\tau) - 3\mu_t(\tau)e^{r\tau} V_t(\tau) + 2\mu_t(\tau)^3}{[e^{r\tau} V_t(\tau) - \mu_t(\tau)^2]^{3/2}} \\
KURT_t^Q(\tau) &= \frac{e^{r\tau} X_t(\tau) - 4\mu_t(\tau)e^{r\tau} W_t(\tau) + 6\mu_t(\tau)^2 e^{r\tau} V_t(\tau) - 3\mu_t(\tau)^4}{[e^{r\tau} V_t(\tau) - \mu_t(\tau)^2]^2}.
\end{aligned}$$

where $\mu_t(\tau) \equiv \mathbb{E}^Q[R(t, \tau)] \approx e^{r\tau} [1 - \frac{1}{2}V_t(\tau) - \frac{1}{6}W_t(\tau) - \frac{1}{24}X_t(\tau)] - 1$.

The computed RNMs are those of the underlying futures return $\ln(\frac{F_T}{F_t})$. Note that the S&P 500 futures price converges to the underlying spot S&P 500 index at maturity, and present futures price is equal to the current spot index price plus cost of carry, ignoring futures margins. The cost of carry is essentially risk-free rate less aggregate index dividends carried over short time intervals of 10-minnutes. For the period 2009 to 2012, the low risk-free rate of less than 0.5% per annum and aggregate S&P 500 dividend yield of 2% implies a basis point negative cost of carry of about 0.04 basis points over short 10-minutes interval. This is an extremely small value that likely implies the risk-neutral moments we compute on the underlying index futures are similar to the risk-neutral moments on underlying index prices. Thus the implications of this study on the S&P 500 index futures prices would be similar to implications on the S&P 500 index prices.

2.2 Data and Implementations

The widely disseminated intraday volatility index VIX and also its traded futures price is an indication of changes in market perception of volatility and also sensitivity or fear of market return losses. The level of VIX had been high since 2009 till mid-2012 during the global financial crisis, and then settled to a calmer level ranging between 10 and 20 points up till end 2014. There is therefore a priori a higher chance of observing market anomalies during the more turbulent times of 2009 till 2012. Moreover, to compare with the other cited studies such as Neumann and Skiadopoulos (2013), Amaya, Christoffersen, Jacobs, and Vasquez (2015), Conrad, Dittmar, and Ghysels (2013), and Bali and Murray (2013) which employed data up to or before 2013, we use intra-day E-mini S&P500 European-style options time-stamped traded price data on weekly index futures series (EW1, EW2 and EW4) from August 2009 to December 2012 purchased from the Chicago Mercantile Exchange (CME). The start date of August 2009 was due to the inability of obtaining clean time-stamped data of sufficient frequency on the stated contracts. The stated options typically trade actively two weeks before their expiration dates. The price data are actual transactions data which are time-stamped to the second. They are not mid-points of bid and ask prices used in other studies. The price data reflects a transaction occurring when a buyer decides to move a spread above to transact at an ask price, or else a transaction occurring when a seller decides to move a spread below to transact at a bid price. However, the data available do not provide information on the bid-ask prices nor the volume transacted. Transactions data were similarly obtained on the E-mini futures prices.

A 2013 World Federation of Exchanges Derivatives Market Survey Report and other trading websites show that the S&P 500 E-mini futures and options are some of the most actively traded index futures and futures options in the world. The E-mini futures has an annual trading volume of 57 million contracts compared to 7 million for the standard S&P 500 index options, although the contract size in the latter is 5 times larger. Indeed the index ETF and index options with larger contract values are known to be used more by position players, whereas the E-mini index futures and its options with smaller contract sizes are used more by speculators and short-term hedgers, and are more liquid as their time to maturity are short, with highest liquidity in the last ten days of trading.

We only consider E-mini options (futures options) traded between 0830hrs and 1500hrs that correspond to regular trading hours, and ignore options traded on the expiration date itself. This is because options are generally illiquid during irregular trading hours, and abnormal option prices tend to occur more frequently on the day of expiration. We incorporate all options, including in-the-money (ITM) options, to capture the information

contained in these options. As the underlying asset is the S&P 500 index futures, we also obtain the E-mini S&P 500 Futures intra-day price data from CME. The options data are cleaned by removing a small percentage of prices that violated no-arbitrage bounds discussed in Merton (1973). The risk-free rate used in our computation is the yield on the secondary market 4-weeks Treasury bills reported in the Federal Reserve Report H.15. Since the risk-free rates in our sample period are close to zero, we can safely assume the risk-free rate is constant throughout any given trading day without incurring non-trivial approximation error.

In order to calculate the risk-neutral moments via the BKM method, we need to compute $V(t, \tau)$, $W(t, \tau)$, and $X(t, \tau)$ that require in theory an infinite number of OTM call and put option prices across a continuum of strike prices. In reality, however, market option prices are available for only a finite number of discrete strike prices even in highly liquid markets. See Jiang and Tian (2005) for a discussion of how to estimate the RNM by joining the call and put price points in a smooth manner. In general, the more call and put option prices are available, the less is any bias due to the discrete number of option prices. As all call and put prices in $V(t, \tau)$, $W(t, \tau)$, and $X(t, \tau)$ are no-arbitrage prices, thus any call or put prices with same maturity formed out of other no-arbitrage call and put prices at the same time t must also be part of $V(t, \tau)$, $W(t, \tau)$, and $X(t, \tau)$. To increase the number of OTM options for the purpose of computing $V(t, \tau)$, $W(t, \tau)$, and $X(t, \tau)$, we consider in-the-money (ITM) call and put options as well to augment mere use of OTM options in standard practice. These ITM options do carry information about their underlying asset returns and should not be unnecessarily ignored.

To incorporate these ITM options, once we apply the data filters, we use the put-call parity equation to transform ITM call option prices into OTM put option prices, and ITM put option prices into OTM call option prices. In this way we are able to increase the number of option price observations in our sample space by including these synthetic OTM option prices and thereby reducing the computation bias of $V(t, \tau)$, $W(t, \tau)$, and $X(t, \tau)$ considerably. Since the put-call parity equation is a model-free no-arbitrage formula, by using this transformation, we are able to utilize the information content in these ITM options and at the same time use it to calculate the risk-neutral moments without imposing any model bias. We should mention that in our empirical results we did try the smaller data sample without the above augmentation, and the results are approximately similar though weaker in the case with smaller sample size.

For each 10-minutes time interval within the regular trading hours, we use as many strike prices as are possible where options were traded on those strike prices. In order to study the ability of market to absorb information within a very short time interval, we

prefer 10-minutes to 30-minutes and so on; we find that choosing interval shorter than 10-minutes is not feasible as there would be some intervals with insufficient number of traded option prices to extract the risk-neutral moments. Also, we find that in many cases 15-minute intervals do not yield any significant change in the reported results. For traded calls (puts) having the same strike price within each 10-minutes band, we select only the call (put) price that was transacted closest to the end of the 10-minutes interval. These call and put prices are then used for computing the risk-neutral moments at every intraday 10-minutes interval starting at 8:40 am, 8:50 am, . . . , 2:40 pm, up to 2:50 pm.

For moment extraction in the intraday intervals in our sample, we use at least two OTM calls and two OTM puts. After that, we employ the numerical method of piece-wise cubic Hermite interpolation to evaluate the integrals for finding the risk-neutral moments. Piece-wise cubic hermite interpolation has a local smoothing property, and therefore produces more stable estimates as compared to cubic splines. The extracted 1 day to 10 day constant maturity risk-neutral moments are used subsequently in our analysis. We do not use options with maturities longer than 10 days because trading for longer E-mini options are less liquid and the price data are inadequate for the purpose of constructing the risk-neutral moments. The results of extracting these risk-neutral moments are reported in Table I below.

Table I about here

Table I reports the descriptive statistics of the extracted S&P 500 risk-neutral moments including the risk-neutral volatility, risk-neutral skewness, and risk-neutral kurtosis. E-mini S& P 500 options with different time-to-maturity of 1 day to 10 days are used to produce the risk-neutral moments corresponding to the different time-to-maturity of n days. Altogether 19,859 moments, one for each 10-minutes interval, are estimated. For the risk-neutral volatility, each n -day volatility is scaled by $\sqrt{252/n}$ so that they are easily compared on an annual basis. The volatilities are reported in %. The risk-neutral skewness and kurtosis, however, are reported without any scaling as these quantities do not have simple distribution-free aggregation properties. Except for volatility which is reported in %, the other moments are reported in decimals.

The averages of moments across all 10-minutes intervals on all dates where their maturities are the same n -day are reflected as the mean for the n -day. Table I shows that the mean volatility is quite stable across all maturities. The mean and median skewness are all negative for all maturities. This is consistent with similar results reported in the literature. Kurtosis generally decreases as maturities increase. Most of the kurtosis measures are in excess of 3, indicating large deviations from the normal distribution. While the

skewness measures are left-skewed, the volatility and kurtosis measures are right-skewed in their frequency distributions over time.

Our results indicate that the risk-neutral distributions on average have more negative skewness and much higher kurtosis than that of normal distributions. For each horizon n , the risk-neutral moments computed for each 10-minutes intervals are highly variable as can be seen by their standard deviations and maximum-minimum ranges. This may reflect the large amounts of new information impacting the market on 10-minutes intervals.

The annualized standard deviations of the various 10-minutes interval risk-neutral moments show that the risk-neutral volatility (Panel A) is the most variable and the intraday time series process is not smooth. On the other hand, the risk-neutral skewness (Panel B) shows a smoother process as its risk-neutral moments evolve slowly through any day. In our sample, 97.9% of all intraday risk-neutral skewness measures are negative. The risk-neutral kurtosis (Panel C) is also relatively smooth given its annualized standard deviations are also smaller than those of the risk-neutral volatility. The smoother process may imply higher chances of successful forecasting.

When we compare the descriptive statistics with the 1-month, 2-month, and 3-month risk-neutral moments reported in Neumann and Skiadopoulou (2013) Table 1, it is seen that annualized volatilities for different horizons remain rather constant for even up to 3 months. Similarly, risk-neutral skewness remains on average negative and does not appear to change much. However, it is clearly the case that our 1 to 10 day risk-neutral kurtosis measures are larger than the longer horizon 30-day to 90-day kurtosis measures.

3 Forecasting Models

In this section we describe the regression models that we use to fit the time series of each of the risk-neutral moments. For each of the 3 different risk-neutral moments, the regressions are performed also on different time series belonging to the different constant maturities. For ease of exposition, we use the notation $RNM_t(\tau)$ to generically represent any of the 3 risk-neutral moments at time t , or the t^{th} interval, with maturity τ . Also $RNV_t(\tau)$, $RNS_t(\tau)$, and $RNK_t(\tau)$ denote risk-neutral volatility, risk-neutral skewness, and risk-neutral kurtosis respectively. Each time series $RNM_t(\tau)$ is associated with a particular trading day. Thus for each trading day, each τ , and each risk-neutral moment, we run a forecasting model.

Firstly we perform a test of the stationarity of each of the risk-neutral moment time series. Without taking up more space, we report that for all cases of maturities, each t , and each RNM, the series are stationary based on the augmented Dickey-Fuller tests. In

all cases, unit root is rejected at less than 1% significance level. This is different from the results in Neumann et. al. where daily risk-neutral moments can display unit roots, so that they require use of ARIMA and ARIMAX processes for forecasting. We do not need to use forecasting models on integrated processes.

We consider 8 competing models in this paper: a benchmark Random Walk Model (RW), an Autoregressive (AR) lag-one Model, an Autoregressive Moving Average Model (ARMA(1,1)), an Autoregressive (AR(1)) Model with GARCH (generalized autoregressive conditional heteroskedastic) error – AR(G), Vector Autoregressive (VAR) lag-one Model where all three lagged RNMs enter as regressors, Vector Autoregressive (VAR) lag-one Model with GARCH errors for each of the three vector elements – VAR(G), Vector Error Correction Model (VECM), and the Local Autoregressive (LAR) lag-one Model. Experimenting with higher lag-orders generally does not yield any clearer results or improvement in analyses. As the lag-order is understood, we shall not clutter the notation and leave out the lag-one notation. In what follows, each interval $[t, t + 1)$ is 10-minutes within a trading day.

For the RW Model, for each RNM:

$$RNM_{t+1}(\tau) = RNM_t(\tau) + \epsilon_{t+1} ,$$

where ϵ_{t+1} is an i.i.d. noise.

For the AR Model, for each RNM:

$$RNM_{t+1}(\tau) = b_0 + b_1 RNM_t(\tau) + \epsilon_{t+1} ,$$

where b_0 and $b_1 < 1$ are constants and ϵ_{t+1} is i.i.d.

For the ARMA Model, for each RNM:

$$RNM_{t+1}(\tau) = b_0 + b_1 RNM_t(\tau) + \epsilon_{t+1} ,$$

where b_0 and $b_1 < 1$ are constants and ϵ_{t+1} is MA(1), with $\epsilon_{t+1} = \alpha\epsilon_t + \varepsilon_{t+1}$, $\alpha < 1$, and ε_t i.i.d.

For the AR(G) Model, for each RNM:

$$RNM_{t+1}(\tau) = b_0 + b_1 RNM_t(\tau) + \epsilon_{t+1} ,$$

where b_0 and $b_1 < 1$ are constants and $\text{var}(\epsilon_{t+1}) = \alpha_0 + \alpha_1 \text{var}(\epsilon_t) + \alpha_2 \epsilon_{t+1}^2$ with constants $\alpha_0 > 0$, and $\alpha_1 + \alpha_2 < 1$.

For the VAR Model:

$$\begin{pmatrix} RNV_{t+1}(\tau) \\ RNS_{t+1}(\tau) \\ RNK_{t+1}(\tau) \end{pmatrix} = B_0 + B_1 \begin{pmatrix} RNV_t(\tau) \\ RNS_t(\tau) \\ RNK_t(\tau) \end{pmatrix} + e_{t+1},$$

where B_0 is a 3×1 vector of constants, B_1 is a 3×3 matrix of constants, and e_{t+1} is a 3×1 vector of i.i.d. disturbance terms.

For the VAR(G) Model: the above VAR Model is used except that each element of the vector error e_{t+1} is modelled as GARCH (1,1).

For the VECM Model:

$$\begin{pmatrix} \Delta RNV_{t+1}(\tau) \\ \Delta RNS_{t+1}(\tau) \\ \Delta RNK_{t+1}(\tau) \end{pmatrix} = \Gamma_0 + \Gamma_1 \begin{pmatrix} RNV_t(\tau) \\ RNS_t(\tau) \\ RNK_t(\tau) \end{pmatrix} + \Gamma_2 \begin{pmatrix} \Delta RNV_t(\tau) \\ \Delta RNS_t(\tau) \\ \Delta RNK_t(\tau) \end{pmatrix} + e_{t+1},$$

where Γ_0 is a 3×1 vector of constants, Γ_1 and Γ_2 are 3×3 constant matrices, and e_{t+1} is a 3×1 vector of i.i.d. disturbance terms.

For the LAR Model:

$$RNM_{t+1}(\tau) = b_{0,I_d} + b_{1,I_d}RNM_t(\tau) + \epsilon_{t+1},$$

where I_d denotes a subset of the sample points on day d . In the current context, this subset consists of sample data from the latest time point before forecasting, i.e. 12:00 noon, to a lagged time point not earlier than 8:40 am. The statistical procedure in which this subset I_d is selected is explained in the next subsection. LAR basically selects an optimal local window to perform the regression fitting where structural breaks do not occur. While it has the advantage of providing a better fit and possibly better forecast in time series that are not smooth and that may have breaks, the disadvantage is that if the time series is not smooth, the shorter sampling window may yield forecasts and estimates with larger standard errors.

Maximum likelihood regression method, equivalent to least squares in cases of normal random errors, is utilized, except that in the LAR case, the selection of window adds to the regression procedures. Next we explore some intraday dynamics of the risk-neutral moments of the S&P 500 index futures returns or rate of change. The in-sample dynamics are examined using autoregressive regressions. Table II reports summary results of the intraday regressions involving the AR(1) model. It reports in-sample regressions of the risk-neutral moments. For each maturity of 1 to 10 days, the extracted 10-minute interval risk-neutral moments (RNMs) of E-mini S&P 500 futures returns or rate of change over

the sample period August 24, 2009 to December 31, 2012 are used for a regression each trading day. On day d for maturity τ , the following autoregressive AR(1) regression is performed for each of the risk-neutral moments:

$$RNM_{t+1}(\tau) = b_0 + b_1 RNM_t(\tau) + \epsilon_{t+1}$$

where maturity is τ , t denotes the particular 10-minute interval during the trading day d , and ϵ_{t+1} is assumed to be i.i.d.

The estimated \hat{b}_0 and \hat{b}_1 are recorded for every t and across all t in the sample period, their averages are reported as Const and Slope respectively. Note that the averages Const and Slope are reported for each of risk-neutral volatility, skewness, kurtosis, and each maturity from 1 day to 10 days. Next, in-sample, for each day d , for each t within d , a next 10-minute interval forecast of the $RNM_{t+1}(\tau)$ is made using $E_t(RNM_{t+1}(\tau)) = \hat{b}_0 + \hat{b}_1 RNM_t(\tau)$. The sign is noted and compared with actual $RNM_t(\tau)$ change at t , $RNM_{t+1}(\tau) - RNM_t(\tau)$.

If both signs are the same, i.e. correct directional in-sample 'prediction', then correct prediction count is increased by one. Over all such t 's in the sample period, for each τ and each RNM, the % of correct prediction counts is denoted as MCP (in %) and reported in the table. In each RNM, each τ , for each d , each t , if the directional in-sample 'prediction' is correct, conditional on this, the actual magnitude of % $RNM_t(\tau)$ change at t , $\text{Abs}(RNM_{t+1}(\tau)/RNM_t(\tau) - 1)$ is noted. This is averaged across all its occurrences and reported as 'RNM Δ ' in the table.

Table II about here

For risk-neutral volatility and all maturities especially for the shorter ones, the estimated constants and slopes are mostly ($> 50\%$) significantly different from zero at the 10% significance level.¹ For risk-neutral skewness and kurtosis, about 25% estimates are significantly different from zero. RNV regressions have smaller constants of regressions relative to those of the skewness and kurtosis regressions. They also have higher slopes. The RNV regressions generally increase in constants and reduce in slopes as maturity increases. For skewness and kurtosis, the constants and slopes remain about the same with different maturities, although there is a slight decrease in constants for kurtosis. The risk neutral skewness regressions mostly have a negative constant.

Clearly the autoregressive or slope coefficients are estimated to be positive for all RNMs. The lag effects do not appear to diminish with maturity for skewness and kurtosis, but is a reduction in the magnitude of the lag effect for volatility as maturity lengthens.

¹ To save print space, some of the details are not reported here. They are available from the authors.

There is higher persistence in RNV. The significance of the regression coefficient estimates in these in-sample regressions indicates the plausibility of using lagged moments to forecast next 10-minutes moments within the day. The out-of-sample forecasting is done and reported in Section 4 while the use of the latter forecasts for options trading strategies and profitability is done in Section 5.

Besides understanding the characteristics of the intraday RNM dynamics, Table II also provides an indication of how often an autoregressive model could predict correctly the next interval increase or decrease in the actual RNM. Thus the in-sample % of correct prediction counts MCP (in %) is reported in the table. It is seen that within-sample, despite the higher persistence of RNV, RNS and RNK autoregressions effectively produce higher MCPs. The constants, slope of regressions as well as how closely the data bunched up and in fact how closely they follow the regression lines, not just the latter, contribute to the MCP counts. Conditional on a correct prediction, the actual magnitude of % $RNM_t(\tau)$ change at t , $\text{Abs}(RNM_{t+1}(\tau)/RNM_t(\tau) - 1)$ is also larger for RNS and RNK. Thus we have reasons to believe that if out-of-sample processes are similar to within-sample processes, then RNS particularly with higher MCP and $RNM\Delta$ could outperform RNV and RNK in terms of forecasts and also trading profits.

As we saw in Table I, the time series of risk-neutral volatility is less smooth with indication of plausible breaks in the stationarity at some points within the trading day. One method to address this issue with intraday study is to employ local autoregression (LAR) model, a technique that is gaining popularity in statistical analyses but not familiar in financial applications. We employ LAR as an alternative model in our forecasts. In the following subsections, we explain and discuss the implementation of the LAR Model.

3.1 Local autoregressive model

We provide in this subsection a concise discussion of the local autoregressive (AR) model used in our forecast. For each one-dimensional risk neutral moment RNM_t representing any of the 3 risk-neutral moments (risk-neutral volatility RNV_t , risk-neutral skewness RNS_t , and risk-neutral kurtosis RNK_t), the idea is to fit an AR(1) model for out-of-sample forecasting, but by using a sample set of 10-minutes RNMs, I_d on day d , starting at time $t \geq 8 : 40am$ and ending at 12:00 noon.

The optimal I_d is selected from competing sample sets $(t, 12:00 pm)$ such that maximum $t = t'$ provides for 10 observation points in the sample. The question is of course how to identify the sub-sample of local homogeneity. In any day of forecast, we seek the longest sub-sample, beyond which there is a high possibility of structural change occurring. A sequential testing procedure is proposed to detect the optimal sub-sample among a

number of candidates with increasing sample lengths starting from $(t', 12:00 \text{ pm})$. The time dependent estimation windows distinguish the LAR model from the conventional AR model with predetermined sample length. See Chen, Wolfgang, and Pigorsch (2010) for justification of this local adaptive method over traditional time series modeling.

The LAR model is particularly useful when the series of RNMs for estimation and forecasting may not be stationary and contains structural breaks at some time points during intraday trading as is common with breaking news or disturbances due to swift changes in high-frequency trading volumes and directions, and also the switches between dark pools and exchange trading orders. The RNM process is specified as the following.

$$RNM_{t+1} = b_{0,I_n} + b_{1,I_n}RNM_t + \varepsilon_{t+1,I_n}, \quad \varepsilon_{t+1} \sim (0, \sigma_{I_n}^2)$$

where I_n denotes a time interval of the local sub-sample $(t_n, 12:00 \text{ pm})$. Increasing sub-samples are denoted by $I_K \supset I_{K-1} \supset \dots \supset I_1$ where I_1 corresponds to $(t', 12:00 \text{ pm})$.

The local maximum likelihood estimators (MLE) $\tilde{\theta}_n = (\tilde{b}_{0,I_n}, \tilde{b}_{1,I_n}, \tilde{\sigma}_{I_n})^\top$ are obtained via:

$$\begin{aligned} \tilde{\theta}_n &= \arg \max_{\theta \in \Theta} L(RNM; I_n, \theta) \\ &= \arg \max_{\theta \in \Theta} \left\{ -(m_n + 1) \log \sigma_{I_n} - \frac{1}{2\sigma_{I_n}^2} \sum_{j \in I_n} (RNM_{j+1} - b_{0,I_n} - b_{1,I_n}RNM_j)^2 \right\} \end{aligned}$$

where m_n is number of data points in I_n , Θ is the parameter space and $L(RNM; I_n, \theta)$ is the local log-likelihood function.

We start from the shortest sub-sample I_1 , let an adaptive estimator be $\hat{\theta}_1 = \tilde{\theta}_1$. The selection procedure then iteratively extends the sub-sample with more 10-minutes sample points and sequentially tests for possible structural breaks in the next longer sub-sample. The significance of structural breaks is measured by a sequential of log-likelihood ratio tests. Once break is detected before reaching I_K , say at n' , then $\tilde{\theta}_{n'}$ is taken as the optimizer LAR estimator for day d . The test statistic is defined in each iterative step following sub-sample I_n as

$$T_{n+1} = |L(I_{n+1}, \tilde{\theta}_{n+1}) - L(I_n, \tilde{\theta}_n)|^{1/2}, \quad n = 1, \dots, K - 1 \quad (1)$$

where $L(I_n, \tilde{\theta}_n) = \max_{\theta \in \Theta} L(RNM; I_n, \theta)$ denotes fitted log-likelihood. If there does not exist a significant structural break in the time within $I_{n+1} - I_n$, the MLE $\tilde{\theta}_{n+1}$ is not far from the estimate $\tilde{\theta}_n$, and the test-statistic T_{n+1} is small. In this case, the extended sub-sample with more information I_{n+1} is accepted and the corresponding MLE $\tilde{\theta}_{n+1}$ replaces $\tilde{\theta}_n$ for improved estimation accuracy. On the other hand, suppose the test statistic is

significantly large, then the selection procedure terminates and the latest accepted sub-sample I_n yields the optimal MLE $\tilde{\theta}_n$. A set of critical values ζ_1, \dots, ζ_K is calibrated in Monte Carlo experiments and used to determine the significance level of T_n for each n . The details of this calibration can be obtained from the authors.

4 Forecasting Performance

After the regression models are estimated, the estimated coefficients are used to provide a fitted model for the purpose of predicting the next period or future risk-neutral moments. Parameters are estimated in the window on the same day from 8:40 am to 12:00 noon, after which the fitted model is used for forecasting during 12:10 pm to 2:50 pm. Unlike daily or weekly methods we do not use rolling windows over the 10-minutes intervals within a trading day. This helps in focusing on trading days with highly liquid transactions at start of day trading to fix the parameters for forecast and trading for the rest of the day. As mentioned before, some days whereby there are insufficient risk-neutral moments for estimation during 8:40 am to 12:00 noon are excluded. The forecast for the various models are shown as follows. Forecasts are made for risk-neutral moments pertaining to different horizons τ of one up to ten days as described earlier.

For the RW Model:

$$E_{t+1}(RNM_{t+2}(\tau)) = RNM_{t+1}(\tau),$$

where the subscript to the expectation operator denotes a condition on the information of $RNM_{t+1}(\tau)$ at $t + 1$.

For the AR Model, for each RNM:

$$E_{t+1}(RNM_{t+2}(\tau)) = \hat{b}_0 + \hat{b}_1 RNM_{t+1}(\tau),$$

where \hat{b}_0 and \hat{b}_1 are estimated parameters in I_t .

For the ARMA Model, for each RNM:

$$E_{t+1}(RNM_{t+2}(\tau)) = \hat{b}_0 + \hat{b}_1 RNM_{t+1}(\tau) + \hat{\alpha} \hat{\epsilon}_{t+1},$$

where $\hat{\epsilon}_{t+1} = RNM_{t+1}(\tau) - \hat{b}_0 - \hat{b}_1 RNM_t(\tau)$.

For the AR(G) Model, for each RNM:

$$E_{t+1}(RNM_{t+2}(\tau)) = \hat{b}_0^G + \hat{b}_1^G RNM_{t+1}(\tau),$$

where \hat{b}_0^G and \hat{b}_1^G are estimated parameters in I_t based on maximum likelihood procedures

recognizing the GARCH variance in the residuals.

For the VAR Model, for all three RNMs at once:

$$E_{t+1} \begin{pmatrix} RNV_{t+2}(\tau) \\ RNS_{t+2}(\tau) \\ RNK_{t+2}(\tau) \end{pmatrix} = \hat{B}_0 + \hat{B}_1 \begin{pmatrix} RNV_{t+1}(\tau) \\ RNS_{t+1}(\tau) \\ RNK_{t+1}(\tau) \end{pmatrix},$$

where \hat{B}_0 and \hat{B}_1 are estimated parameters in I_t .

For the VAR(G) Model, for all three RNMs at once:

$$E_{t+1} \begin{pmatrix} RNV_{t+2}(\tau) \\ RNS_{t+2}(\tau) \\ RNK_{t+2}(\tau) \end{pmatrix} = \hat{B}_0^G + \hat{B}_1^G \begin{pmatrix} RNV_{t+1}(\tau) \\ RNS_{t+1}(\tau) \\ RNK_{t+1}(\tau) \end{pmatrix},$$

where \hat{B}_0^G and \hat{B}_1^G are estimated parameters in I_t based on GARCH(1,1) errors e_{t+2} .

For the VECM Model, for all three RNMs at once:

$$E_{t+1} \begin{pmatrix} RNV_{t+2}(\tau) \\ RNS_{t+2}(\tau) \\ RNK_{t+2}(\tau) \end{pmatrix} = \hat{\Gamma}_0 + \left(\hat{\Gamma}_1 + I \right) \begin{pmatrix} RNV_{t+1}(\tau) \\ RNS_{t+1}(\tau) \\ RNK_{t+1}(\tau) \end{pmatrix} + \hat{\Gamma}_2 \begin{pmatrix} \Delta RNV_{t+1}(\tau) \\ \Delta RNS_{t+1}(\tau) \\ \Delta RNK_{t+1}(\tau) \end{pmatrix},$$

where $\hat{\Gamma}_0$, $\hat{\Gamma}_1$ and $\hat{\Gamma}_2$ are estimated parameters in I_t

For the LAR Model:

$$E_{t+1} (RNM_{t+2}(\tau)) = \hat{b}_{0,I_{t+1}} + \hat{b}_{1,I_{t+1}} RNM_{t+1}(\tau),$$

where $\hat{b}_{0,I_{t+1}}$ and $\hat{b}_{1,I_{t+1}}$ are the estimated parameters in I_{t+1} .

4.1 Error Metrics

To measure the forecasting performances of these models, we employ 3 error metrics or loss functions.

The Root Mean Square Error (RMSE) is defined as

$$RMSE = \sqrt{\frac{1}{T-1} \sum_{t=0}^{T-2} \left(RNM_{t+2}(\tau) - E_{t+1} (RNM_{t+2}(\tau)) \right)^2},$$

where T is the number of periods of forecasts, each period being a 10-minutes interval.

The Mean Absolute Deviation (MAD) is defined as

$$MAD = \frac{1}{T-1} \sum_{t=0}^{T-2} \left| RNM_{t+2}(\tau) - E_{t+1}(RNM_{t+2}(\tau)) \right|.$$

The Mean Correct Prediction (MCP) percentage is defined as

$$MCP = \frac{1}{T-1} \sum_{t=0}^{T-2} J_{t+2} \times 100,$$

where indicator $J_{t+2} = 1$ if $(RNM_{t+2}(\tau) - RNM_{t+1}(\tau))(E_{t+1}(RNM_{t+2}(\tau)) - RNM_{t+1}(\tau)) > 0$, and $J_{t+2} = 0$ otherwise.

Table III about here

Table III reports the out-of-sample statistical performances of the Random Walk (RW) model, the Autoregressive (AR) model, the Autoregressive Moving Average (ARMA) model, the Autoregressive with Garch error (AR(G)) model, the Vector Autoregressive (VAR) model, the Vector Autoregressive with Garch errors (VAR(G)) model, the Vector Error Correction model (VECM), and the Local Autoregressive (LAR) model. The respective autoregressive models are lag-one models. Results are reported for each of the risk-neutral moments of volatility, skewness, and kurtosis. For each risk-neutral moment (RNM) category, the regression results of all maturities are pooled. There is a total of 8534 observations for each RNM regression. Every trade day from August 24, 2009 to December 31, 2012, the risk-neutral moments computed in each 10-minutes intervals from 8:40 am to 12:00 pm (noon) are used to estimate the parameters of each model. The estimated or fitted model is then used to forecast the risk-neutral moments for each 10-minutes interval from 12:10 pm to 14:50 pm. The error metrics or loss functions of root mean square error (RMSE), mean absolute deviation (MAD), and mean correct prediction (MCP) % are shown in the table. The MCP is the % of times that the forecast of directional change in the risk-neutral moment is correct.

The results in Table III provide clear indications of the following. Firstly, the VECM model performs the worst of all; this shows that more complicated models with more than one lag could yield less accurate forecasts. Similarly, VAR model does not perform well as there may be multi-correlation of the RNM's in a finite sample setting. Secondly, across all RNMs, the AR(G) model, the LAR model, and the AR model perform better than the rest in terms of lower RMSE, lower MAD, and higher MCP. For the more volatile RNV processes, LAR appears to perform slightly better in MAD and in MCP. These latter models are all autoregressive in nature. This may not be surprising since the results in

Table II shows the pervasion of autoregression. Remarkably, all methods perform better than the RW in terms of MCP higher than 50%. In summary, there is statistical evidence of a lot of intraday information that can be utilized to successfully make rather accurate predictions of next period risk-neutral moments over short intervals of 10-minutes.

Next we consider the forecasts of the individual risk-neutral moments, and also consider the forecasts in specific time series corresponding to different maturities. The results are reported in tables IV to VI.

Tables IV to VI about here

Table IV reports the statistical performances of the Random Walk (RW), the Autoregressive (AR) model, the Autoregressive Moving Average (ARMA) model, the Autoregressive with Garch error (AR(G)) model, the Vector Autoregressive (VAR) model, the Vector Autoregressive with Garch errors (VAR(G)) model, the Vector Error Correction model (VECM), and the Local Autoregressive (LAR) model in forecasting. The respective autoregressive models are lag-one models. A separate regression is performed for the risk-neutral moments data corresponding to each time-to-maturity. Every trade day from August 24, 2009 to December 31, 2012, the risk-neutral moments computed in each 10-minutes intervals from 8:40 am to 12:00 pm (noon) are used to estimate the parameters in each regression. The estimated or fitted model is then used to forecast the risk-neutral moments for each 10-minutes interval from 12:10 pm to 14:50 pm. The error metrics or loss functions of root mean square error (RMSE), mean absolute deviation (MAD), and mean correct prediction (MCP) % are shown in panels A, B, and C respectively. The MCP is the % of times that the forecast of directional change in the risk-neutral moment is correct. Note: The risk-neutral volatility in our study is expressed in %. If expressed as decimals, the magnitudes of the RMSE and MAD error metrics would be reduced by a multiple of 10^{-2} .

In Table IV, clearly VECM, VAR, and VAR(G) are mostly underperforming with larger RMSE, MAD, and smaller MCP relative to the other forecasting models. As observed in Table III, AR(G), LAR, and AR or ARMA are better performers in forecasting. All models outperform RW in terms of MCP, though not necessarily in RMSE and MAD. For RMSE and MAD, forecasting performances for all models deteriorate with longer maturities. The MCPs however stay approximately the same for all maturities. For this RNV, LAR yields the best MCP amongst all models except for the longest maturity of 10 days.

In Table V, for skewness, forecasting with AR outperforms with AR(G) and LAR coming in as close seconds. This is the case for all three error metrics. Except for VECM and VAR(G), all models beat the RW model in skewness forecasting on metrics

RMSE and MAD. However, all models outperform the RW in terms of MCP. Distinctly, unlike the RNV, there is no apparent deterioration in forecasting for skewness with longer maturities. In most cases, in fact, the forecasting errors reduce for longer maturities. This suggests that RNMs equations for volatility contain a much larger residual noise, so that apparently higher first order correlations in volatility shown in Table II do not imply better forecasts over longer horizons. All the models beat the RW in terms of MCP. More interestingly, if we examine the MCPs between RNV and RNS forecasting, forecasting skewness is shown to produce more accuracies. This is shown in aggregated form in Table III where there is an average MCP of over 70% for AR, AR(G), and LAR models in skewness forecasting, but only 58% to 69% in the volatility forecasting. In Table V, this phenomenon is uniform across all maturities. Thus again, the Table II indications of higher autocorrelation in RNV than in RNS must not be taken as an interpretation of better performances in prediction for RNV.

In Table VI, for kurtosis, forecasting with LAR outperforms with AR(G) and AR coming in as close seconds. This is the case for all three error metrics. As with RNV and RNS, VECM, VAR(G), and also ARMA models do not perform well relative to the others. Distinctly, unlike the RNV and RNS, the forecast for RNK improves with maturity. Across all models, RMSE reduces from a magnitude of 4-5 in 1-day maturity to a magnitude of 2+ in 10-days maturity. Similarly MAD reduces from 2-3 in 1-day maturity to 1+ in 10 days maturity. The MCPs remain steady at about 69% across the maturities, excluding the case of VECM. Thus, for RNK, the longer the horizon, there appears to be some mean reversion and thus better forecasting performance. This is consistent with the Table I finding that over longer maturities, kurtosis shrinks on average. If we examine the MCPs between RNK, RNV, and RNS forecasting, forecasting kurtosis is shown to produce accuracies in between the cases of RNV and RNS. This is shown in aggregated form in Table III where there is average MCP of about 68% for LAR, AR(G), and AR models in kurtosis forecasting, but only 58% to 69% in the volatility forecasting, and over 70% in skewness forecasting. In Table VI, this phenomenon is uniform across all maturities.

5 Options Trading Strategies

Using the forecasts generated by the 7 competing models of AR lag-one, ARMA(1,1), AR(1) with GARCH error, VAR lag-one, VAR(1) with GARCH errors, VECM, and LAR, we attempt to construct a trading strategy to benefit from accurate forecast of the various future moment changes. RW is excluded as it has served its purpose for benchmark comparison in the forecast assessments shown in Tables III to VI. Besides our 7 competing

models outperform RW in terms of MCP. We also add the benchmark case of perfect knowledge forecast (PK) whereby prediction of moment increase or decrease is 100% correct.

We construct 3 different trading strategies corresponding to the forecasts of the 3 different risk-neutral moments. The trading strategies are aimed for empirical moments forecasts of which the risk-neutral moments forecasts for the look-ahead horizon are supposedly proxies. Although the direction and not the magnitude of the risk-neutral moment prediction is key to the options portfolio strategy, we also incorporate some consideration of the magnitude of the forecast changes – this is called threshold. We consider different thresholds for different RNMs as some like RNS and RNK have forecast ranges that can be much larger compared to RNV. In Table I, standard deviation divided by mean averages about 0.4 for RNV, 0.6 for RNS, and 0.65 for RNK, indicating the kind of ranges. If a threshold is $x\%$, we will execute the options portfolio strategy at a particular 10-minutes interval $t + 2$ when the forecast of the RNM exceeds the RNM in the last interval by at least $x\%$, i.e. $E_{t+1}(RNM_{t+2}(\tau))/RNM_{t+1} \geq 1 + x\%$.

Also, unlike other studies that do not explicitly account for the cost of transaction for each trade, i.e. commission cost, our option trading strategy needs to weigh the cost of a trade versus the ex-ante signal of gain provided by the forecast of future moment changes. If we do not set a threshold > 0 , then any forecast signal no matter how small, would trigger a trade, and this could be very costly. On the other hand, if we set a threshold to be too high, then most forecast moments changes would not trigger a trade, and there would be too few trade and the sample of trade profits would be too small and carry large standard errors.

In practice, using thresholds based on expected forecast changes is much more parsimonious and clearer to the trader than trying to use statistical significance levels on the parameters of each regressions or using other forecast metrics as decision thresholds to trade or not. However, setting thresholds cannot be arbitrary from the point of view of ex-ante research – based on examination of the sampling data from the first couple of months, suitable thresholds yielding enough trades and producing positive profits are chosen for the rest of the sampling period. Based on a range of such suitable thresholds, profitability results are reported as follows – some of these results are not necessarily significant from an ex-post point of view.

For trading, the portfolio involves typically some E-mini calls, puts, and also the E-mini futures contracts. With our dataset of actual transacted prices in calls, puts, and futures, we are able to micmic closely actual trading possibilities. We assume that in the transaction at 10-minutes time interval $t + 2$, from the time at $t + 1$, we are able to

trade with market orders at those transacted prices at $t + 2$. In our execution, we use only actual traded option prices and do not use synthetic prices that we had created for the purpose of only developing the forecasting equations in the morning part of the day. In the variance and skewness cases we have enough of options typically, though for the kurtosis case, as it involves four options within a 10-minutes interval, there were some days where there were signals but no trades due to shortages of transactions prices. This shows up in Table IX where there are less numbers of trades. We also assume trading in a small number of portfolio units each time interval in order not to impact on the spread. Our results show per round trip portfolio profit per unit, so the results are valid regardless of the trading impact.

For forecast on risk-neutral volatility, the trading strategy involves creating a volatility portfolio each 10-minutes interval as follows: long an OTM call and short delta amount of underlying asset, together with long an OTM put and short a delta amount of underlying index futures. The respective deltas are based on the strike prices of the call and the put of the same maturity, and are computed using Black-Scholes model. Since the delta of a put is negative, shorting delta related to a put amounts to buying the underlying index futures asset. Each trading interval is 10-minutes intraday. At the end of the interval the positions are liquidated at actual market prices. Prediction is done on moments with the same maturity. If the predicted next interval risk-neutral volatility is higher than the current risk-neutral volatility by at least the threshold percentage, the above portfolio of long call and long put is executed.

The execution and liquidation next interval constitute one trade. There can be more than one trade per interval if different maturity moment forecasts exceed the threshold, or there may be no trade in a particular interval if a non-zero threshold signal is used. The portfolio has zero cost as the net balance of the cost in the call, put, and underlying asset is financed by borrowing at risk-free rate. In the opposite case when predicted risk-neutral volatility is lower than the current risk-neutral volatility by at least the threshold percentage, the portfolio is short. In this case, a call and a put are sold together with associated positions in the underlying asset and risk-free bond. The overall cost of the portfolio can be expressed as:

$$\pi_t^{Vola} = C_{t,OTM} - \Delta_{C_{t,OTM}} F_t + P_{t,OTM} - \Delta_{P_{t,OTM}} F_t - B_t.$$

B_t is chosen such that $\pi_t^{Vola} = 0$. Outlay for F_t the index futures is assumed to be zero for the initial futures position. The profit measure can be evaluated as a change in the above cost. We impute a trading commission cost of 22.5 cents per option contract or 45 cents per round trade of an option contract position. This is an average figure obtained

from brokers dealing with large orders or familiar clients.² Trading takes place from 12:10 pm till 15:00 pm each trading day based on the forecasts.

Table VII about here

Table VII reports the average trading profit in \$ per round-trip trade net of commission costs according to the different forecasting methods on RNV and the various threshold signals of 0.0%, 5.0%, 7.5%, and 10.0%. As the theoretical probability distribution of the average profit is not known, and the sample size is not very large, we employ the bootstrap method suggested by Efron (1979) to compute the standard errors for the statistics. The numbers in the brackets indicate the t-statistics based on bootstrapped variances calculated for the average profit. The bootstrap is carried out over 2000 iterations of the sample size indicated by the total number of trades. **, * denote significance at the 1-tail 1%, and 2.5% significance levels respectively. The number of trades #Trades refers to all possible trades where there are signals given enough options during 8:40 am till 12:00 noon for making forecast and where there are enough actual calls and puts during 12:10 pm till 14:50 pm for execution each 10-minutes period. #Trades for AR(Garch) and VAR(Garch) are less because of some cases of non-convergence in the estimation algorithm. Generally higher signals produced less number of trades because there are less signals for trade execution.

Table VII shows that for 0% and 5% thresholds, trading the volatility induce losses. If we add back the commission costs that were deducted, most of the models yield a before-cost \$ payoff of close to zero per trade. Thus greater than 50% MCPs are not a guarantee for profitable options strategy on volatility. The forecast is on the next 10-minutes expected risk-neutral volatility while the options strategy is to capture empirical volatility on the underlying futures, so the strategy may not capture profit if the difference between risk-neutral and empirical volatility is significant. However, when the threshold is 7.5%, large variance change signals produce in LAR a \$1.33 average profit, even though the t-statistic is not significant even at 5% level. Overall, risk-neutral volatility cannot lead to profitable options trading strategy. This is similar to results using daily risk-neutral moments as in Neumann et.al.

² We find that most studies do not indicate the size of their trading commission costs, being part of transactions costs. Therefore we venture that what we have is an estimate. For example, the internet trading firm "TradeStation" offers flat-fee trading of \$4.99 + \$0.20 per option trade for trading volumes of 200 or more contracts per month. This works out to about \$ 0.20 + \$4.99/200 = \$ 0.225 per option trade. For a large institutional broker with a seat on the relevant exchange, it is plausible for this level of low commission cost. The futures position costs would even be smaller for institutional brokers and are close to zero for very large contract sizes.

For forecast on skewness, the trading strategy involves creating a skewness portfolio each 10-minutes interval as follows: long an OTM call and short a number of OTM puts equal to the ratio of the call vega to put vega. Also short a number of underlying assets equal to the call delta less the same vega ratio times put delta. The respective vegas and deltas are based on the strike prices and other features of the call and the put. The vegas and deltas are computed based on Black-Scholes model. Each trading interval is 10-minutes intraday. At the end of the interval the positions are liquidated at market prices. If the predicted next interval risk-neutral skewness is higher than the current risk-neutral skewness by at least the threshold percentage, the above portfolio of long call and short put of the same maturity is executed. The execution and liquidation next interval constitute one round-trip trade. There can be more than one trade per interval if different maturity moment forecasts exceed the threshold, or there may be no trade in a particular interval if a non-zero threshold signal is used. The portfolio has zero cost as the net balance of the cost in the call, puts, and underlying asset is financed by borrowing at risk-free rate. In the opposite case when predicted risk-neutral skewness is lower than the current risk-neutral skewness by at least the threshold percentage, the portfolio is short. In this case, there is a short call and a long put together with associated positions in the underlying asset and risk-free bond. The overall cost of the portfolio can be expressed as:

$$\pi_t^{Skew} = C_{t,OTM} - \left(\frac{v_{C_{t,OTM}}}{v_{P_{t,OTM}}} \right) P_{t,OTM} - \left(\Delta_{C_{t,OTM}} - \left(\frac{v_{C_{t,OTM}}}{v_{P_{t,OTM}}} \right) \Delta_{P_{t,OTM}} \right) F_t - B_t.$$

B_t is chosen such that $\pi_t^{Skew} = 0$. Outlay for F_t the index futures is assumed to be zero for the initial futures position. The profit measure can be evaluated as a change in the above cost. Similarly we impute a trading commission cost of 22.5 cents per option contract or 45 cents per round trade of an option contract position.

Table VIII about here

Table VIII reports the average trading profit in \$ per round-trip trade net of commission costs according to the different forecasting methods on RNS and the various threshold signals of 0.0%, 10.0%, 20.0%, and 50.0%. As the theoretical probability distribution of the average profit is not known, and the sample size is not very large, we employ the bootstrap method suggested by Efron (1979) to compute the standard errors for the statistics. The numbers in the brackets indicate the t-statistics based on bootstrapped variances calculated for the average profit. The bootstrap is carried out over 2000 iterations of the sample size indicated by the total number of trades.

In the skewness case as shown in Table VIII, except for VECM forecasting, all the other forecasting models using thresholds less than 50% produce significantly positive

\$ profits per trade at 1% significance level. The AR and AR(G) methods consistently produce high \$ average trading profit per trade of \$1.40 to \$2 for thresholds from 0% up to 20%. Since the results are obtained using transactions prices and not mid of bid-ask prices, the profits already accounted for spread cost. Moreover, trading commissions costs are also deducted. Thus, the intraday 10-minutes forecasts and trades produce significant profits not seen in the daily trading results elsewhere such as in Neumann et.al.

The possibility of risk-neutral skewness forecast on index futures return to yield profitable trading strategies using options in the next 10-minutes interval indicates that the E-mini options market is not informationally efficient. This could be the reason why this market sees very active high frequency trading by speculators as well as hedgers. Informational inefficiency is construed in the sense of one being able to use available information at $t+1$ to profit from trades at $t+2$. In so far as the risk-neutral skewness is close to empirical skewness, the options portfolio strategy is one which benefits from the subsequent manifestation of index futures returns skewness. This skewness profitability anomaly may be an indication of informational market inefficiency in intraday S&P 500 futures options markets. This is different from inefficiency in the index futures markets since most of the trades in a portfolio execution are positions in the futures options. If positive skewness forecast results in long position in OTM calls and short position in OTM puts, being able to profit implies that the calls are underpriced relative to the puts or the puts are overpriced relative to the calls. This in itself does not however imply that there is breach to no-arbitrage conditions on the option prices. An example to illustrate could be that an OTM call on strike 1400 is valued using a no-arbitrage model at 1.935, with underlying index futures at 1300. An OTM put with strike 1200, same maturity, is priced at 1.353. One can make profit by buying the OTM call, selling the OTM put, hedging the vega and delta, if the skewness indeed increases with index futures rising to 1350. However, the calls and puts themselves are non-arbitrageable and their prices follow put-call parities.

For forecast of kurtosis, the trading strategy involves creating a kurtosis portfolio each 10-minutes interval as follows: long X number of ATM calls and X number of ATM puts and simultaneously short one OTM call and one OTM put, where $X = (C_{t,OTM} + P_{t,OTM}) / (C_{t,ATM} + P_{t,ATM})$. Each trading interval is 10-minutes intraday. At the end of the interval the positions are liquidated at market prices. Prediction is done on moments with the same maturity. If the predicted next interval risk-neutral kurtosis is higher than the current risk-neutral kurtosis by at least the threshold percentage, the above portfolio of long ATM call and put, and short OTM call and put, all of the same maturity, is executed. The execution and liquidation next interval constitute one trade. There can be more than one trade per interval if different maturity moment forecasts exceed the

threshold, or there may be no trade in a particular interval. The portfolio is self-financing and has zero cost.

In the opposite case when predicted risk-neutral kurtosis is lower than the current risk-neutral kurtosis by at least the threshold percentage, the portfolio is short. In this case, there is short X numbers of ATM calls and ATM puts, and long one of OTM call and OTM put. The overall cost of the portfolio can be expressed as:

$$X(C_{t,ATM} + P_{t,ATM}) - (C_{t,OTM} + P_{t,OTM}) = 0.$$

Table IX about here

Table IX reports the average trading profit in \$ per round-trip trade net of commission costs according to the different forecasting methods on RNK and the various threshold signals of 0.0%, 10.0%, 20.0%, and 50.0%. For a predicted increase in kurtosis, Ait-Sahalia, Wang, and Yared (2001)'s study on S&P 500 options indicated that the risk-neutral probabilities would increase in prices closer to the ATM strike while they would decrease in prices far OTM. Therefore the trading strategy as indicated above would be to buy ATM or near ATM options and sell OTM options. Table IX shows that there are far fewer trades possible under this kurtosis strategy since four options with same maturity have to be traded in each 10-minutes intervals and there are fewer of such intervals. There are only two cases of significant profits. At 20% threshold, the AR model yields a significant profit at 2.5% level. At 10% threshold, the VECM yields a significant profit at 1% level. However, the VECM model does not appear to be consistently profitable across other thresholds. Only the AR and AR(G) models appear to produce more consistent profitability across the various thresholds, but these are not largely significant.

We also attempt an approximate comparison of our trading strategies with the daily trading profitability reported in Neumann and Skiadopoulos (2013). We qualify the approximation as the data used in our study are different in several aspects, even as both studies are analysing risk-neutral moments of distributions related to the S&P 500 index. Neumann et.al. use S&P 500 index quote and mid-point of bid-ask quoted option prices as the underlying. We use actual transaction prices on the index futures and the index futures options (E-mini futures options). They use end of trading day option prices in OptionMetrics; we use CME intraday transaction prices in the last 10 minutes of the trading hours each trading day. The OptionMetrics data spanned 30 days and so on, while our data for EW options are meant for short maturities of one up to ten days. It is more difficult to find enough transacted option prices within a 10-minutes end-of-trading day span, so our daily data points are more limited. Neumann et.al. remove data where im-

plied volatility exceeds 100%; we do not, as long as all data meet no-arbitrage conditions. This may allow higher volatility and kurtosis to show up on our data.

Unlike use of S&P 500 index, or even its proxy of SPY ETFs which tend to be traded by larger institutional funds on position taking, our E-mini futures is more a day trading vehicle which suits the essence of our study. Our daily data of 506 points starts basically in 2011-2102 as dense option trading at end of day is not available in the earlier period. We use for each day, RNM computed on the longest maturity contracts, and average maturity is just over 5 days. In our daily 1-day ahead forecast, we employ a constant rolling 60 days window of daily computed RNMs. Similar skewness and kurtosis trading strategies based on RNM signals with no threshold i.e. 0%, as in Tables VII, VIII, IX, are employed over 1 day. The Sharpe ratios of the trading results based on daily trading profits are reported as follows. Following Neumann et.al. we perform similar bootstrap and report the 95% confidence in the pair of numbers in the brackets below the averaged Sharpe ratio.

Table X about here

In the case of skewness, as in the the Neumann and Skiadopoulos study, our Sharpe ratios are positive, though not significant at 1-tail 2.5% level. For the case of kurtosis, our forecasting methods yield positive Sharpe ratios in the use of ARMA, AR(G), and VAR(G) models whereas Neumann et.al.'s kurtosis strategy does not appear to perform as well, making mostly losses. We do not attempt to compare daily trading outcomes in risk-neutral volatility as our study does not show intraday profitability based on RNV forecast and similarly Neumann et.al. find that predictability of the risk-neutral moments do not lead to profitable outcomes in volatility trading.

Our results show no profitability net of trading cost in all forecasts of RNMs whereas Neumann et.al. (Table 10) displayed for comparison shows some profitability in skewness forecasts before deducting transaction costs. They are not significantly larger than 0 at 1-tail 2.5% level, but are reported in Nuemann et.al. as significantly larger than 0, before spread costs, at 1-tail 5% level. The conclusions from both studies with respect to daily trading profitability look similar.

The sharp difference of extant studies with ours is that we are testing for intraday options trading whereas almost all significant published studies so far had concentrated on trading over at least one day interval if not weeks or months. The longer kinds of trading intervals are not typical of a trader with a smart strategy and higher trading frequencies in a highly liquid market such as S&P 500 derivatives instruments. Also, there is a marked difference in our options trading strategies, such as in kurtosis, compared to

those for example in the Neumann et. al. study. They used forecast of 60-day and 90-day implied moments compared to our intraday moments; their study also employed delta and vega hedged portfolios to attempt to extract profits from kurtosis changes. For trading intervals longer than a day, such hedging ratios may not be accurate and may contain high sampling risks, so their use are not particularly effective in capturing advantages from a change in kurtosis. On the other hand, we rely on a keen observation reported in Aït-Sahalia, Wang, and Yared (2001) and develop a much cleaner and effective kurtosis strategy as is evident in our comparative trading results.

6 Ex-Ante Moments and Subsequent Returns

In single period asset pricing or intertemporal asset pricing with stationary stochastic investment opportunities, there are a majority of models where higher systematic variance or covariance, systematic co-kurtosis, and lower co-skewness are compensated by higher expected returns. These measures pertain to every security under equilibrium pricing, and are measures related to variance, skewness and kurtosis, except the latter can be diversified in a portfolio but the former cannot. Theoretical papers include Kraus and Litzenberger (1976) and Harvey and Siddique (2000) on co-skewness, and Dittmar (2002) on co-kurtosis.

In a somewhat similar setup, portfolios of stocks that are left-skewed, having high variance, and high kurtosis, can achieve a certain degree of within portfolio diversification (though not the general equilibrium model diversification), and thus across these portfolios, their average returns will display higher ex-post return compensation. This is empirically shown in the works of Bali and Murray (2013) and Conrad, Dittmar, and Ghysels (2013). A number of other papers have presented some evidence of ex-ante skewness being negatively related to stock returns. This relationships is about equilibrium asset pricing and does not directly have anything to do with informational market efficiency. However, they are consistent with the possibility that for each security, the contemporaneous return is related to its other moments; if idiosyncratic noise for a stock is small, then clearly this linear relationship will allow for the moments to be significantly priced (as in factor pricing models). Such a linear representation of return on its own risk-neutral moments is of course an approximate proxy for an equilibrium model. In our case we study the S&P 500 index futures or equivalently the index, and there is so far scant research on the predictability of index return itself based on past moments. Ratcliff (2013) is one study that documents daily ex-ante risk-neutral skewness innovations having a negative impact on S&P 500 returns.

If we could forecast RNMs as in Table II or use more sophisticated forecasting models in Tables IV to VI, then we could examine the possibility also of forecasting returns in the next period. In our context we consider forecasting over the next short 10-minutes interval. As actual returns are highly variable, it would be challenging to find predictability on futures returns itself. For the forecast regression we may write:

$$\ln\left(\frac{F_{t+1}(\tau)}{F_t(\tau)}\right) = \beta_0 + \beta_1 \ln\left(\frac{F_t(\tau)}{F_{t-1}(\tau)}\right) + \beta_2 \left(E_t[RNM_{t+1}(\tau)] - RNM_t(\tau)\right) + \epsilon_{t+1}(\tau),$$

where $t = 1, \dots, T$ are the different 10-minutes intervals in any trading day, $\epsilon_{t+1}(\tau)$ is the i.i.d. residual error, τ indicates the correspondence of the futures price and RNM to futures and futures options with maturity τ , and the second explanatory variable is the ex-ante RNM innovation. This innovation is based on an AR(1) forecast of $RNM_{t+1}(\tau)$. Note we add the lagged futures return in order to capture any bid-ask bounce trades.

A separate forecasting regression involving all three RNM's as explanatory variables is also run: $\ln\left(\frac{F_{t+1}}{F_t}\right) * 100 = \beta_0 + \beta_1 \ln\left(\frac{F_t}{F_{t-1}}\right) * 100 + \beta_2 \left(E_t[RNV_{t+1}(\tau)] - RNV_t(\tau)\right) + \beta_3 \left(E_t[RNS_{t+1}(\tau)] - RNS_t(\tau)\right) + \beta_4 \left(E_t[RNK_{t+1}(\tau)] - RNK_t(\tau)\right) + \epsilon_{t+1}(\tau)$.

Table XI about here

In Table XI, we display results for $\tau = 1, 5$ days. (Results for other τ 's are basically similar.) For each trading day where there are many trades throughout the day, the regression is performed and the coefficient estimates are collected. Over the total number of such trading days, the estimated coefficients are averaged to obtain their means. Their t-statistics, within brackets reflecting if their average deviates significantly from zero, are also evaluated and reported in this table. The results for the above regressions are reported in each column for different models 1 to 4. Models 1,2, and 3 use RNV,RNS, and RNK respectively. Model 4 uses all the RNMs as explanatory variables. Bootstrap is carried out over 2000 iterations of the sample size indicated by the total number of trades. **, *, + denote 2-tail 5%, 2%, 1% significance levels respectively.

The results show that as in other studies there is a statistically significant negative bid-ask bounce appearing in the estimated coefficients of the lagged futures returns. The risk-neutral volatility contributes negatively to next interval futures returns, though the estimated coefficients are far from significant with low t-statistics based on bootstrapped distribution. Ex-ante risk-neutral skewness innovation contributes negatively to next interval returns for most maturities, though again the estimated coefficients are not significant. This seems to be consistent with daily results concerning negative relationship

between ex-ante skewness and future returns of stocks, though the latter is cross-sectional across stocks and not a time series result. It is different from the significant daily regression results of Ratcliff (2013). Ex-ante risk-neutral kurtosis have close to zero effect and is insignificant, like all the other ex-ante RNMs. Thus the futures market itself, unlike the options market, appears to pass the test of efficiency in intraday trading using moments information.

7 Conclusions

As far as we know, our paper is one of the first studies on intraday implied moments of S&P 500 index futures returns using intraday or high frequency futures option prices and the E-mini index futures prices. The E-mini S&P500 options are European-style and the data are actual transactions prices on the CME Exchange. We improve on existing techniques by using cubic hermite interpolation with better smoothing properties to extract the first four moments of the risk-neutral return distribution. Secondly we perform intraday out-of-sample forecasting or prediction, and document the intraday dynamics of the index futures return risk-neutral moments. We also introduce a novel local autoregression method that allows variable window in fitting the autoregressive parameters. This is particularly useful in situations when there may be intraday news that cause structural changes in the returns or price distributions.

Based on the forecasts of the risk-neutral moments on each 10-minutes interval during a trading day, option portfolios are constructed, executed for trade and then liquidated 10-minutes later to take advantage of the forecast information. The trades are performed on market orders and trading commission costs are also charged. The key results we find are that there is profitability after transactions costs in the trading strategies involving risk-neutral skewness, and that autoregressive models including those modeling GARCH errors on the moments innovations are the most accurate in terms of less mean absolute errors and higher directional accuracy in forecasting. An improved kurtosis strategy recognizing the empirical observation of higher risk-neutral densities around ATM relative to OTM option prices when forecast kurtosis increases leads to very marginal cases of profitable trading based on kurtosis forecasts. The last result on kurtosis is quite interesting as past studies typically found trading on kurtosis to be losses.

The positive profitability after transaction costs in skewness trading indicates an anomaly which may be an indication of market inefficiency in the intraday S&P 500 futures options markets – it could be due to information inefficiency within short spans of time. We thus provide fresh results in intraday trading contrary to the findings of

informational efficiency in daily trading by studies in the period up to 2013 which is also approximately the period of our study. This also provides an explanation of the observed phenomenon that many practitioners, such as institutional traders, continue to trade intraday for profits when research had maintained that the market is informationally efficient.

Finally we also attempt to examine if our autoregressive models of ex-ante risk-neutral moments forecasts could help predict the next 10-minutes underlying index futures returns. We find the expected bid-ask bounce effects, but no predictive ability by the expected moments innovations. This is different from the Ratcliff (2013) study using daily data. Thus the futures market itself, unlike the options market, appears to pass the test of efficiency in intraday trading using moments information. This is not surprising given that the underlying S&P 500 futures market has extremely high liquidity. This last result is not contrary to other related studies that find a negative asset pricing equilibrium relationship between ex-ante skewness and stock returns in cross sectional analyses.

Table I
Descriptive Statistics of S&P 500 Risk-Neutral Moments

This table reports the descriptive statistics of the extracted S&P 500 risk-neutral moments including the risk-neutral volatility, risk-neutral skewness, and risk-neutral kurtosis in different panels A, B, and C respectively. E-mini S&P 500 futures options with different time-to-maturities of 1 day to 10 days are used to produce the risk-neutral moments corresponding to the different time-to-maturities of n days. The options data are collected over the sample period of August 24, 2009 to December 31, 2012. Altogether 19,859 10-minutes intervals of estimates of risk-neutral moments are obtained. For the risk-neutral volatility, each n -day volatility is scaled by $\sqrt{252/n}$ so that they are easily compared on an annual basis. The risk-neutral volatility numbers are computed in %. The risk-neutral skewness and kurtosis, however, are reported without any scaling as these quantities do not have simple distribution-free aggregation properties.

Panel A: Risk-Neutral Volatility (%)										
Time-to-Maturity n	1 Day	2 Days	3 Days	4 Days	5 Days	6 Days	7 Days	8 Days	9 Days	10 Days
No. of Obs.	2628	2769	2594	2277	1920	1987	1784	1471	1205	1224
Mean	21.05	19.85	18.97	19.07	20.49	20.31	20.35	20.89	20.43	21.41
Median	18.47	17.30	17.02	17.10	17.64	17.60	18.29	18.14	17.60	17.73
Max	67.44	85.90	54.26	84.68	81.59	84.20	72.56	79.24	83.25	77.49
Min	9.70	9.35	8.84	8.83	8.61	8.92	9.42	9.61	9.65	7.93
Std Dev.(%)	8.21	7.68	6.90	7.32	7.91	8.18	7.34	8.25	8.43	8.97
Skewness	1.85	2.14	1.62	2.329	1.98	1.83	1.41	1.38	1.61	1.32
Kurtosis	7.29	10.19	5.78	12.27	10.12	8.51	6.31	6.86	6.10	4.62

Panel B: Risk-Neutral Skewness										
Time-to-Maturity n	1 Day	2 Days	3 Days	4 Days	5 Days	6 Days	7 Days	8 Days	9 Days	10 Days
No. of Obs.	2628	2769	2594	2277	1920	1987	1784	1471	1205	1224
Mean	-1.01	-1.03	-1.07	-1.17	-1.26	-1.12	-1.08	-1.05	-1.03	-1.02
Median	-0.95	-0.97	-1.00	-1.10	-1.21	-1.08	-1.05	-1.04	-0.97	-1.02
Max	2.34	2.07	1.03	1.80	2.26	0.87	0.88	1.02	0.88	1.18
Min	-3.65	-3.81	-3.53	-3.17	-3.45	-3.14	-3.00	-3.64	-3.68	-3.27
Std Dev.(%)	0.62	0.58	0.61	0.61	0.62	0.52	0.48	0.51	0.52	0.56
Skewness	-0.41	-0.59	-0.60	-0.46	-0.31	-0.23	-0.22	-0.04	-0.61	0.15
Kurtosis	3.89	4.79	4.00	3.92	4.36	4.01	4.21	4.16	5.19	4.61

Panel C: Risk-Neutral Kurtosis										
Time-to-Maturity n	1 Day	2 Days	3 Days	4 Days	5 Days	6 Days	7 Days	8 Days	9 Days	10 Days
No. of Obs	2628	2769	2594	2277	1920	1987	1784	1471	1205	1224
Mean	6.28	5.72	5.76	6.28	6.53	5.16	4.68	4.49	4.37	4.14
Median	4.90	4.56	4.48	4.98	5.38	4.34	4.06	3.81	3.58	3.59
Max	30.11	29.60	28.97	30.11	29.70	24.70	19.47	21.25	28.91	26.40
Min	0.88	0.01	0.03	0.02	0.02	0.02	0.03	0.03	0.03	0.03
Std Dev.(%)	4.41	3.94	4.00	4.20	4.30	3.09	2.60	2.62	2.97	2.64
Skewness	2.11	2.32	2.09	1.77	1.89	1.90	1.87	1.83	2.89	2.50
Kurtosis	8.73	10.32	8.54	6.69	7.41	8.30	8.09	7.95	16.21	14.17

Table II
Statistics of In-Sample Regressions

This table reports in-sample regressions of the risk-neutral moments. For the risk-neutral volatility, each n -day volatility is scaled by $\sqrt{252/n}$ so that they are easily compared on an annual basis. The risk-neutral volatility numbers are computed in decimals. The risk-neutral skewness and kurtosis, however, are reported without any scaling as these quantities do not have simple distribution-free aggregation properties. For each maturity of 1 to 10 days, the extracted 10-minute interval risk-neutral moments (RNMs) of E-mini S&P 500 futures returns or rate of change over the sample period August 24, 2009 to December 31, 2012 are used for a regression each trading day. On day d for maturity τ , the following autoregressive AR(1) regression is performed for each of the risk-neutral moments:

$$RNM_{t+1}(\tau) = b_0 + b_1 RNM_t(\tau) + \epsilon_{t+1}$$

where maturity is τ , t denotes the particular 10-minute interval during the trading day d , and ϵ_{t+1} is assumed to be i.i.d. The estimated \hat{b}_0 and \hat{b}_1 are recorded for every t and across all t in the sample period, their averages are reported as Const and Slope respectively. Note that the averages Const and Slope are reported for each of risk-neutral volatility, skewness, kurtosis, and each maturity from 1 day to 10 days. Next, in-sample, for each day d , for each t within d , a next 10-minute interval forecast of the $RNM_{t+1}(\tau)$ is made using $E_t(RNM_{t+1}(\tau)) = \hat{b}_0 + \hat{b}_1 RNM_t(\tau)$. The sign is noted and compared with actual $RNM_t(\tau)$ change at t , $RNM_{t+1}(\tau) - RNM_t(\tau)$. If both signs are the same, i.e. correct directional in-sample 'prediction', then correct prediction count is increased by one. Over all such t 's in the sample period, for each τ and each RNM, the % of correct prediction counts is denoted as MCP (in %) and reported in the table. In each RNM, each τ , for each d , each t , if the directional in-sample 'prediction' is correct, conditional on this, the actual magnitude of % $RNM_t(\tau)$ change at t , $\text{Abs}(RNM_{t+1}(\tau)/RNM_t(\tau) - 1)$ is noted. This is averaged across all its occurrences and reported as 'RNM Δ ' in the table.

τ	Risk-Neutral Volatility				Risk-Neutral Skewness				Risk-Neutral Kurtosis			
	Const	Slope	MCP	RNM Δ	Const	Slope	MCP	RNM Δ	Const	Slope	MCP	RNM Δ
1 Day	0.077	0.60	63.7	0.03	-0.922	0.13	72.0	1.13	5.388	0.16	70.0	0.54
2 Days	0.082	0.52	62.9	0.04	-0.932	0.11	72.7	1.22	4.972	0.16	70.9	0.70
3 Days	0.093	0.45	65.4	0.04	-0.919	0.08	72.5	1.14	4.786	0.14	71.5	0.62
4 Days	0.106	0.42	64.2	0.05	-0.997	0.16	69.9	1.88	5.324	0.15	70.7	1.49
5 Days	0.124	0.35	63.3	0.06	-1.115	0.12	73.5	1.11	5.616	0.16	71.1	2.00
6 Days	0.127	0.32	64.3	0.05	-0.971	0.06	72.8	1.01	4.263	0.11	71.1	1.21
7 Days	0.137	0.27	65.5	0.05	-0.969	0.05	72.6	0.73	3.989	0.12	72.6	0.86
8 Days	0.152	0.27	67.5	0.06	-0.924	0.09	71.8	2.74	3.635	0.19	72.0	0.95
9 Days	0.153	0.24	67.5	0.07	-0.890	0.15	70.7	0.73	3.440	0.23	68.9	0.73
10 Days	0.169	0.18	69.0	0.06	-0.929	0.12	70.8	1.28	3.643	0.15	70.1	0.65

Table III
Out-of-Sample Error Metrics for Forecasting Models

This table reports the out-of-sample performances of the Random Walk (RW) model, the Autoregressive (AR) model, the Autoregressive Moving Average (ARMA) model, the Autoregressive with Garch error (AR(G)) model, the Vector Autoregressive (VAR) model, the Vector Autoregressive with Garch errors (VAR(G)) model, the Vector Error Correction model (VECM), and the Local Autoregressive (LAR) model. The respective autoregressive models are lag-one models. Results are reported for each of the risk-neutral moments of volatility, skewness, and kurtosis. For each risk-neutral moment (RNM) category, the regression results of all maturities are pooled. There is a total of 8534 observations for each RNM regression. Every trade day from August 24, 2009 to December 31, 2012, the risk-neutral moments computed in each 10-minutes intervals from 8:40 am to 12:00 pm (noon) are used to estimate the parameters of each model. The estimated or fitted model is then used to forecast the risk-neutral moments for each 10-minutes interval from 12:10 pm to 14:50 pm. The error metrics or loss functions of root mean square error (RMSE), mean absolute deviation (MAD), and mean correct prediction (MCP) % are shown in the table. The MCP is the % of times that the forecast of directional change in the risk-neutral moment is correct.

	Risk-Neutral Volatility			Risk-Neutral Skewness			Risk-Neutral Kurtosis		
	RMSE	MAD	MCP	RMSE	MAD	MCP	RMSE	MAD	MCP
RW	0.662	0.174	50.00%	0.647	0.463	50.00%	4.001	2.490	50.00%
AR	0.580	0.175	58.15%	0.511	0.374	70.75%	3.252	2.174	68.98%
ARMA	0.598	0.175	59.66%	0.581	0.424	69.44%	3.619	2.373	68.10%
AR(G)	0.498	0.162	61.18%	0.521	0.378	70.27%	3.230	2.131	69.63%
VAR	0.654	0.210	58.35%	0.541	0.387	69.87%	3.465	2.255	68.40%
VAR(G)	0.569	0.182	63.04%	0.625	0.433	67.83%	3.387	2.209	68.27%
VECM	3.619	0.396	60.36%	1.358	0.544	63.20%	7.995	3.092	61.99%
LAR	0.526	0.159	68.67%	0.535	0.374	71.01%	3.391	2.110	70.23%

Table IV

Out-of-Sample Error Metrics for Forecasting Risk-Neutral Volatility

This table reports the statistical performances of the Random Walk (RW), the Autoregressive (AR) model, the Autoregressive Moving Average (ARMA) model, the Autoregressive with Garch error (AR(G)) model, the Vector Autoregressive (VAR) model, the Vector Autoregressive with Garch errors (VAR(G)) model, the Vector Error Correction model (VECM), and the Local Autoregressive (LAR) model. The respective autoregressive models are lag-one models. A separate regression is performed for moments data corresponding to each time-to-maturity. Every trade day from August 24, 2009 to December 31, 2012, the risk-neutral moments computed in each 10-minutes intervals from 8:40 am to 12:00 pm (noon) are used to estimate the parameters in each regression. The estimated or fitted model is then used to forecast the risk-neutral moments for each 10-minutes interval from 12:10 pm to 14:50 pm. The error metrics or loss functions of root mean square error (RMSE), mean absolute deviation (MAD), and mean correct prediction (MCP) % are shown in panels A, B, and C respectively. The MCP is the % of times that the forecast of directional change in the risk-neutral moment is correct. Note: The risk-neutral volatility in our study is expressed in %. If expressed as decimals, the magnitudes of the RMSE and MAD error metrics would be reduced by a multiple of 10^{-2} .

Panel A: Root-Mean-Square Error (RMSE)										
Time-to-Maturity n	1 Day	2 Days	3 Days	4 Days	5 Days	6 Days	7 Days	8 Days	9 Days	10 Days
No. of Obs.	1151	1225	1141	993	814	844	759	629	498	480
RW	0.079	0.296	0.258	0.827	0.996	0.499	0.779	0.904	0.900	0.998
AR	0.078	0.286	0.193	0.679	0.911	0.454	0.680	0.771	0.920	0.747
ARMA	0.086	0.263	0.195	0.750	0.923	0.520	0.689	0.735	0.924	0.806
AR(G)	0.079	0.226	0.208	0.639	0.743	0.448	0.563	0.590	0.691	0.757
VAR	0.138	0.294	0.213	0.660	1.222	0.466	0.650	0.962	0.925	0.809
VAR(G)	0.128	0.301	0.423	0.663	0.825	0.526	0.611	0.703	0.717	0.853
VECM	0.085	1.735	1.247	3.525	8.384	1.499	1.311	3.549	7.066	1.283
LAR	0.065	0.206	0.176	0.652	0.785	0.498	0.541	0.520	0.933	0.798

Panel B: Mean Absolute Deviation (MAD)										
Time-to-Maturity n	1 Day	2 Days	3 Days	4 Days	5 Days	6 Days	7 Days	8 Days	9 Days	10 Days
No. of Obs.	1151	1225	1141	993	814	844	759	629	498	480
RW	0.056	0.089	0.094	0.195	0.244	0.169	0.240	0.274	0.273	0.382
AR	0.057	0.097	0.092	0.187	0.251	0.200	0.243	0.270	0.279	0.307
ARMA	0.060	0.097	0.093	0.189	0.237	0.192	0.243	0.267	0.290	0.342
AR(G)	0.057	0.090	0.094	0.174	0.217	0.196	0.212	0.235	0.247	0.315
VAR	0.105	0.134	0.126	0.219	0.311	0.225	0.260	0.301	0.322	0.329
VAR(G)	0.063	0.105	0.117	0.189	0.257	0.206	0.235	0.268	0.270	0.352
VECM	0.061	0.197	0.153	0.496	0.995	0.344	0.400	0.616	0.760	0.483
LAR	0.045	0.078	0.079	0.173	0.226	0.193	0.227	0.232	0.265	0.313

Panel C: Mean Correct Prediction (MCP%)										
Time-to-Maturity n	1 Day	2 Days	3 Days	4 Days	5 Days	6 Days	7 Days	8 Days	9 Days	10 Days
No. of Obs.	1151	1225	1141	993	814	844	759	629	498	480
RW	50.00	50.00	50.00	50.00	50.00	50.00	50.00	50.00	50.00	50.00
AR	52.80	52.98	58.58	56.84	61.49	57.33	59.97	64.03	62.22	66.81
ARMA	58.87	54.49	60.00	56.65	59.16	60.02	60.55	64.33	63.38	68.89
AR(G)	58.53	54.69	60.47	61.57	62.82	59.88	62.60	66.45	66.60	69.56
VAR	54.28	53.64	57.61	57.55	59.13	59.50	61.73	65.00	61.40	62.96
VAR(G)	63.06	59.55	64.67	61.67	62.06	64.36	65.95	62.58	63.73	65.33
VECM	61.65	59.89	62.46	60.69	57.44	60.74	59.89	59.55	59.96	59.29
LAR	71.33	68.93	69.13	67.55	66.96	67.67	66.71	70.16	68.38	68.95

Table V

Out-of-Sample Error Metrics for Forecasting Risk-Neutral Skewness

This table reports the statistical performances of the Random Walk (RW), the Autoregressive (AR) model, the Autoregressive Moving Average (ARMA) model, the Autoregressive with Garch error (AR(G)) model, the Vector Autoregressive (VAR) model, the Vector Autoregressive with Garch errors (VAR(G)) model, the Vector Error Correction model (VECM), and the Local Autoregressive (LAR) model. The respective autoregressive models are lag-one models. A separate regression is performed for moments data corresponding to each time-to-maturity. Every trade day from August 24, 2009 to December 31, 2012, the risk-neutral moments computed in each 10-minutes intervals from 8:40 am to 12:00 pm (noon) are used to estimate the parameters in each regression. The estimated or fitted model is then used to forecast the risk-neutral moments for each 10-minutes interval from 12:10 pm to 14:50 pm. The error metrics or loss functions of root mean square error (RMSE), mean absolute deviation (MAD), and mean correct prediction (MCP) % are shown in panels A, B, and C respectively. The MCP is the % of times that the forecast of directional change in the risk-neutral moment is correct.

Panel A: Root-Mean-Square Error (RMSE)										
Time-to-Maturity n	1 Day	2 Days	3 Days	4 Days	5 Days	6 Days	7 Days	8 Days	9 Days	10 Days
No. of Obs.	1151	1225	1141	993	814	844	759	629	498	480
RW	0.720	0.647	0.628	0.627	0.704	0.594	0.573	0.606	0.641	0.696
AR	0.566	0.508	0.500	0.508	0.551	0.473	0.461	0.479	0.496	0.530
ARMA	0.629	0.577	0.577	0.565	0.649	0.532	0.516	0.564	0.573	0.604
AR(G)	0.574	0.525	0.508	0.526	0.554	0.481	0.463	0.484	0.493	0.568
VAR	0.579	0.519	0.516	0.591	0.629	0.477	0.483	0.503	0.534	0.541
VAR(G)	0.688	0.627	0.656	0.593	0.734	0.575	0.527	0.568	0.574	0.600
VECM	0.742	1.455	1.011	1.559	1.798	0.668	1.582	1.085	2.435	0.829
LAR	0.610	0.529	0.540	0.515	0.592	0.489	0.475	0.492	0.494	0.550

Panel B: Mean Absolute Deviation (MAD)										
Time-to-Maturity n	1 Day	2 Days	3 Days	4 Days	5 Days	6 Days	7 Days	8 Days	9 Days	10 Days
No. of Obs.	1151	1225	1141	993	814	844	759	629	498	480
RW	0.539	0.477	0.456	0.444	0.486	0.421	0.416	0.432	0.419	0.503
AR	0.422	0.377	0.369	0.377	0.395	0.346	0.339	0.355	0.348	0.379
ARMA	0.470	0.428	0.419	0.414	0.455	0.393	0.380	0.408	0.404	0.451
AR(G)	0.428	0.388	0.371	0.385	0.395	0.351	0.340	0.350	0.342	0.394
VAR	0.430	0.384	0.379	0.402	0.420	0.348	0.353	0.364	0.369	0.388
VAR(G)	0.491	0.443	0.437	0.429	0.467	0.405	0.382	0.406	0.395	0.418
VECM	0.539	0.538	0.484	0.559	0.653	0.459	0.557	0.527	0.617	0.570
LAR	0.435	0.376	0.375	0.364	0.395	0.348	0.341	0.348	0.333	0.386

Panel C: Mean Correct Prediction (MCP%)										
Time-to-Maturity n	1 Day	2 Days	3 Days	4 Days	5 Days	6 Days	7 Days	8 Days	9 Days	10 Days
No. of Obs.	1151	1225	1141	993	814	844	759	629	498	480
RW	50.00	50.00	50.00	50.00	50.00	50.00	50.00	50.00	50.00	50.00
AR	71.24	70.41	70.80	69.29	72.55	69.35	71.56	71.29	69.40	72.16
ARMA	69.48	69.94	68.68	68.15	71.22	68.21	70.84	69.91	68.21	70.15
AR(G)	70.87	70.02	70.48	69.73	71.95	68.36	71.31	70.00	68.24	71.46
VAR	70.80	69.17	69.92	68.67	71.43	69.47	69.68	70.16	66.74	73.02
VAR(G)	67.76	66.92	66.28	65.60	69.92	68.00	70.38	66.29	68.03	72.52
VECM	65.13	65.20	63.95	61.90	63.22	61.09	63.46	61.46	60.36	62.84
LAR	70.98	71.16	71.06	71.22	72.05	68.15	72.24	70.32	72.07	71.31

Table VI
Out-of-Sample Error Metrics for Forecasting Risk-Neutral Kurtosis

This table reports the statistical performances of the Random Walk (RW), the Autoregressive (AR) model, the Autoregressive Moving Average (ARMA) model, the Autoregressive with Garch error (AR(G)) model, the Vector Autoregressive (VAR) model, the Vector Autoregressive with Garch errors (VAR(G)) model, the Vector Error Correction model (VECM), and the Local Autoregressive (LAR) model. The respective autoregressive models are lag-one models. A separate regression is performed for moments data corresponding to each time-to-maturity. Every trade day from August 24, 2009 to December 31, 2012, the risk-neutral moments computed in each 10-minutes intervals from 8:40 am to 12:00 pm (noon) are used to estimate the parameters in each regression. The estimated or fitted model is then used to forecast the risk-neutral moments for each 10-minutes interval from 12:10 pm to 14:50 pm. The error metrics or loss functions of root mean square error (RMSE), mean absolute deviation (MAD), and mean correct prediction (MCP) % are shown in panels A, B, and C respectively. The MCP is the % of times that the forecast of directional change in the risk-neutral moment is correct.

Time-to-Maturity n	1 Day	2 Days	3 Days	4 Days	5 Days	6 Days	7 Days	8 Days	9 Days	10 Days
No. of Obs.	1151	1225	1141	993	814	844	759	629	498	480
RW	5.224	4.245	3.987	4.085	4.471	3.432	2.939	2.968	3.521	2.867
AR	4.261	3.444	3.293	3.535	3.594	2.667	2.345	2.400	2.627	2.194
ARMA	4.689	3.875	3.724	3.870	3.895	3.094	2.512	2.657	3.153	2.392
AR(G)	4.221	3.466	3.254	3.558	3.518	2.660	2.365	2.322	2.542	2.228
VAR	4.406	3.541	3.427	4.068	3.997	2.696	2.553	2.553	2.915	2.278
VAR(G)	4.641	3.679	3.356	3.617	3.595	2.692	2.392	2.418	2.678	2.314
VECM	5.343	6.819	6.520	10.591	13.931	3.687	8.889	5.607	9.228	2.973
LAR	4.327	3.731	3.383	3.787	3.690	2.772	2.494	2.405	2.682	2.384

Time-to-Maturity n	1 Day	2 Days	3 Days	4 Days	5 Days	6 Days	7 Days	8 Days	9 Days	10 Days
No. of Obs.	1151	1225	1141	993	814	844	759	629	498	480
RW	3.309	2.677	2.560	2.616	2.813	2.183	1.902	1.926	1.944	1.854
AR	2.817	2.343	2.271	2.431	2.365	1.891	1.672	1.699	1.701	1.511
ARMA	3.082	2.583	2.472	2.605	2.561	2.127	1.762	1.823	1.960	1.646
AR(G)	2.760	2.331	2.200	2.403	2.318	1.857	1.638	1.642	1.633	1.518
VAR	2.856	2.420	2.348	2.594	2.524	1.920	1.752	1.757	1.784	1.542
VAR(G)	2.963	2.433	2.264	2.433	2.359	1.916	1.686	1.676	1.722	1.551
VECM	3.443	3.284	2.930	3.598	4.206	2.472	2.754	2.478	2.749	1.990
LAR	2.739	2.251	2.167	2.403	2.324	1.826	1.656	1.596	1.630	1.531

Time-to-Maturity n	1 Day	2 Days	3 Days	4 Days	5 Days	6 Days	7 Days	8 Days	9 Days	10 Days
No. of Obs.	1151	1225	1141	993	814	844	759	629	498	480
RW	50.00	50.00	50.00	50.00	50.00	50.00	50.00	50.00	50.00	50.00
AR	69.76	70.33	68.60	67.65	70.43	67.79	69.00	69.52	66.32	68.95
ARMA	68.00	68.38	66.93	67.44	69.74	67.30	69.39	69.27	65.79	69.31
AR(G)	70.43	69.26	68.96	69.73	70.18	68.85	71.31	71.29	66.60	68.71
VAR	68.44	68.51	67.72	67.76	69.32	68.39	69.54	68.87	66.53	68.95
VAR(G)	68.74	66.50	67.26	68.39	69.67	67.64	69.97	70.32	65.16	70.40
VECM	63.13	60.95	62.98	61.69	61.87	60.02	63.59	64.97	57.14	62.42
LAR	72.03	70.74	69.39	69.59	70.68	68.75	70.49	71.45	67.76	70.24

Table VII

Profitability of Trading Strategy using Risk-Neutral Volatility Prediction

The table reports the average \$ trading profit per trade according to the different forecasting methods and threshold signals. Trading cost per option contract is 22.5 cents, and this has been deducted to arrive at the net trading profit. The trading strategy involves creating a volatility portfolio each 10-minutes interval as follows: long an OTM call and short delta amount of underlying asset, together with long an OTM put and short a delta amount of underlying index futures. The respective deltas are based on the strike prices of the call and the put, and are computed using Black-Scholes model. Since the delta of a put is negative, shorting delta related to a put amounts to buying the underlying index futuresasset. Each trading interval is 10-minutes intraday. At the end of the interval the positions are liquidated at actual market prices. Prediction is done on moments with the same maturity. If the predicted next interval risk-neutral volatility is higher than the current risk-neutral volatility by at least the threshold percentage, the above portfolio of long call and long put is executed. The execution now and liquidation next interval constitute one trade. There can be more than one trade per interval if different maturity moment forecasts exceed the threshold, or there may be no trade in a particular interval if a non-zero threshold signal is used. The portfolio has zero cost as the net balance of the cost in the call, put, and underlying asset is financed by borrowing at risk-free rate. In the opposite case when predicted risk-neutral volatility is lower than the current risk-neutral volatility by at least the threshold percentage, the portfolio is short. In this case, a call and a put are sold together with associated positions in the underlying asset and risk-free bond. The overall cost of the portfolio can be expressed as:

$$\pi_t^{Vol} = C_{t,OTM} - \Delta_{C_{t,OTM}}F_t + P_{t,OTM} - \Delta_{P_{t,OTM}}F_t - B_t.$$

B_t is chosen such that $\pi_t^{Vol} = 0$. Outlay for F_t the index futures is assumed to be zero for the initial futures position. The numbers in the brackets indicate the t-statistics based on bootstrapped variances calculated for the average profit. The bootstrap is carried out over 2000 iterations of the sample size indicated by the total number of trades. **, * denote significance at the 1-tail 1%, and 2.5% significance levels respectively. The number of trades #Trades refers to all possible trades where there are signals given enough options during 8:40 am till 12:00 noon for making forecast and where there are enough actual calls and puts during 12:10 pm till 14:50 pm for execution each 10-minutes period. #Trades for AR(Garch) and VAR(Garch) are less because of some cases of non-convergence in the estimation algorithm. Generally higher signals produced less number of trades because there are less signals for trade execution.

Threshold	Signal 0.0%		Signal 5.0%		Signal 7.5%		Signal 10.0%	
	Profit	#Trades	Profit	#Trades	Profit	#Trades	Profit	#Trades
PK	1.91** (7.86)	7137	2.56** (4.51)	2152	3.39** (3.98)	1264	4.31** (4.50)	773
AR	-1.01 (-4.14)	7137	-0.32 (-0.52)	1403	0.46 (0.56)	921	0.87 (0.90)	664
ARMA	-0.96 (-3.92)	7137	-1.40 (-2.08)	1175	-0.54 (-0.55)	719	-0.56 (-0.46)	457
AR(G)	-1.06 (-4.14)	6455	-0.90 (-1.34)	1243	-0.21 (-0.23)	806	0.24 (0.24)	570
VAR	-0.71 (-2.88)	7137	-0.61 (-1.66)	3121	-0.05 (-0.09)	1962	0.59 (0.91)	1296
VAR(G)	-0.69 (-2.68)	6455	-0.69 (-1.80)	2817	-0.07 (-0.14)	1754	0.61 (0.90)	1144
VECM	-0.87 (-3.57)	7137	-0.80 (-1.26)	2175	-0.64 (-0.74)	1265	-1.30 (-1.31)	954
LAR	-0.20 (-0.80)	7137	0.09 (0.15)	2050	1.33 (1.55)	1235	0.84 (0.98)	884

Table VIII

Profitability of Trading Strategy using Risk-Neutral Skewness Prediction

The table reports the average \$ trading profit per trade according to the different forecasting methods and threshold signals. Trading cost per option contract is 22.5 cents, and this has been deducted to arrive at the net trading profit. The trading strategy involves creating a skewness portfolio each 10-minutes interval as follows: long an OTM call and short a number of OTM puts equal to the ratio of the call vega to put vega. Also short a number of underlying assets equal to the call delta less the same vega ratio times put delta. The respective vegas and deltas are based on the strike prices and other features of the call and the put and are computed using Black-Scholes model. Each trading interval is 10-minutes intraday. At the end of the interval the positions are liquidated at actual market prices. Prediction is done on moments with the same maturity. If the predicted next interval risk-neutral skewness is higher than the current risk-neutral skewness by at least the threshold percentage, the above portfolio of long call and short puts is executed. The execution and liquidation next interval constitute one trade. There can be more than one trade per interval if different maturity moment forecasts exceed the threshold, or there may be no trade in a particular interval if a non-zero threshold signal is used. The portfolio has zero cost as the net balance of the cost in the call, puts, and underlying asset is financed by borrowing at risk-free rate. In the opposite case when predicted risk-neutral skewness is lower than the current risk-neutral skewness by at least the threshold percentage, the portfolio is short. In this case, there is a short call and a long put together with associated positions in the underlying asset and risk-free bond. The overall cost of the portfolio can be expressed as:

$$\pi_t^{Skew} = C_{t,OTM} - \left(\frac{v_{C_{t,OTM}}}{v_{P_{t,OTM}}}\right)P_{t,OTM} - \left(\Delta_{C_{t,OTM}} - \left(\frac{v_{C_{t,OTM}}}{v_{P_{t,OTM}}}\right)\Delta_{P_{t,OTM}}\right)F_t - B_t.$$

B_t is chosen such that $\pi_t^{Skew} = 0$. Outlay for F_t the index futures is assumed to be zero for the initial futures position. The numbers in the brackets indicate the t-statistics based on bootstrapped variances calculated for the average profit. The bootstrap is carried out over 2000 iterations of the sample size indicated by the total number of trades. **, * denote significance at the 1-tail 1%, and 2.5% significance levels respectively. The number of trades #Trades refers to all possible trades where there are signals given enough options during 8:40 am till 12:00 noon for making forecast and where there are enough calls and puts during 12:10 pm till 14:50 pm for execution each 10-minutes period. #Trades for AR(Garch) and VAR(Garch) are less because of some cases of non-convergence in the estimation algorithm. Generally higher signals produced less number of trades because there are less signals for trade execution.

Threshold	Signal 0.0%		Signal 10.0%		Signal 20.0%		Signal 50.0%	
	Profit	#Trades	Profit	#Trades	Profit	#Trades	Profit	#Trades
PK	2.81** (9.33)	7137	2.97** (8.98)	5657	3.00** (7.96)	4291	2.68** (5.56)	2406
AR	1.42** (4.76)	7137	1.74** (5.01)	5178	1.88** (4.30)	3290	1.16 (1.81)	1507
ARMA	1.39** (4.65)	7137	2.01** (5.04)	4101	1.84** (3.73)	2645	0.52 (0.76)	1384
AR(G)	1.48** (4.66)	6455	1.80** (4.85)	4661	1.97** (4.21)	2928	1.03 (1.46)	1317
VAR	1.59** (5.28)	7137	1.94** (5.58)	5246	1.75** (4.18)	3336	1.02 (1.67)	1562
VAR(G)	1.62** (5.19)	6455	1.91** (5.23)	4710	1.78** (3.96)	2961	0.91 (1.39)	1372
VECM	0.45 (1.50)	7137	0.91** (2.55)	5115	0.53 (1.21)	3439	0.02 (0.03)	1685
LAR	1.40** (4.65)	7137	1.62** (4.67)	5458	1.83** (4.43)	3647	1.06 (1.67)	1622

Table IX

Profitability of Trading Strategy using Risk-Neutral Kurtosis Prediction

The table reports the average \$ trading profit per trade according to the different forecasting methods and threshold signals. Trading cost per option contract is 22.5 cents, and this has been deducted to arrive at the net trading profit. The trading strategy involves creating a kurtosis portfolio each 10-minutes interval as follows: long X ATM calls and X ATM puts and simultaneously short one OTM call and one OTM put, where $X = (C_{t,OTM} + P_{t,OTM}) / (C_{t,ATM} + P_{t,ATM})$. The ATM (OTM) options are chosen as far as possible to have similar strikes. Each trading interval is 10-minutes intraday. At the end of the interval the positions are liquidated at actual market prices. Prediction is done on moments with the same maturity. If the predicted next interval risk-neutral kurtosis is higher than the current risk-neutral kurtosis by at least the threshold percentage, the above portfolio of long ATM calls and puts, and short OTM call and put is executed. The execution and liquidation next interval constitute one trade. There can be more than one trade per interval if different maturity moment forecasts exceed the threshold, or there may be no trade in a particular interval if a non-zero threshold signal is used. The portfolio has zero cost. In the opposite case when predicted risk-neutral kurtosis is lower than the current risk-neutral kurtosis by at least the threshold percentage, the portfolio is short. In this case, there are short ATM calls and puts, and long OTM call and put. The overall cost of the portfolio can be expressed as:

$$X(C_{t,ATM} + P_{t,ATM}) - (C_{t,OTM} + P_{t,OTM}) = 0.$$

The numbers in the brackets indicate the t-statistics based on bootstrapped variances calculated for the average profit. The bootstrap is carried out over 2000 iterations of the sample size indicated by the total number of trades. **, * denote significance at the 1-tail 1%, and 2.5% significance levels respectively. The number of trades #Trades refers to all possible trades where there are signals given enough options during 8:40 am till 12:00 noon for making forecast and where there are enough calls and puts during 12:10 pm till 14:50 pm for execution each 10-minutes period. #Trades for the Kurtosis strategy are smaller as there are fewer 10-minutes intervals during 12:00 to 14:50 pm whereby actual sets of more options were to be found, i.e. both ATM and OTM calls and puts. Generally higher signals produced less number of trades because there are less signals for trade execution.

Threshold	Signal 0.0%		Signal 10.0%		Signal 20.0%		Signal 50.0%	
	Profit	#Trades	Profit	#Trades	Profit	#Trades	Profit	#Trades
PK	2.92** (2.98)	682	2.70* (2.01)	377	3.53 (1.86)	228	4.35 (1.15)	87
AR	0.71 (0.71)	682	2.82 (1.72)	319	5.78* (2.09)	137	4.00 (0.63)	33
ARMA	0.55 (0.55)	682	2.85 (1.44)	241	1.99 (0.68)	128	5.33 (0.82)	41
AR(G)	0.60 (0.56)	623	2.57 (1.43)	286	5.72 (1.85)	121	4.64 (0.63)	29
VAR	1.37 (1.39)	682	2.16 (1.33)	328	4.42 (1.52)	142	4.18 (0.74)	38
VAR(G)	1.36 (1.28)	623	2.22 (1.27)	301	3.93 (1.25)	129	4.52 (0.76)	36
VECM	0.79 (0.79)	682	3.82** (2.69)	374	3.75 (1.75)	192	1.10 (0.32)	60
LAR	-0.89 (-0.88)	682	1.64 (1.17)	374	3.65 (1.65)	188	2.25 (0.38)	45

Table X
Approximate Comparison of Daily Trading Profitability

An approximate comparison is attempted with the daily trading profitability reported in Neumann and Skiadopoulos (2013). We qualify the approximation as the data used in our study are different in several aspects, even as both studies are analysing risk-neutral moments of distributions related to the S&P 500 index. Neumann et.al. use S&P 500 index quote and mid-point of bid-ask quoted option prices as the underlying. We use actual transaction prices on the index futures and the index futures options (E-mini futures options). Neumann et.al. use end of trading day option prices in OptionMetrics; we use CME intraday transaction prices in the last 10 minutes of the trading hours each trading day. The OptionMetrics data spanned 30 days and so on, while our data for EW options are meant for short maturities of one up to ten days. It is more difficult to find enough transacted option prices within a 10-minutes end-of-trading day span, so our daily data points are more limited. Neumann et.al. removes data where implied volatility exceeds 100%; we do not, as long as all data meet no-arbitrage conditions. This may allow higher volatility and kurtosis to show up on our data. Unlike use of S&P 500 index, or even its proxy of SPY ETFs which tend to be traded by larger institutional funds on position taking, our E-mini futures is more a day trading vehicle which suits the essence of our study. Our daily data of 506 points starts basically in 2011-2102 as dense option trading at end of day is not available in the earlier period. We use for each day, RNM computed on the longest maturity contracts, and average maturity is just over 5 days. In our daily 1-day ahead forecast, we employ a constant rolling 60 days window of daily computed RNMs. Similar skewness and kurtosis trading strategies based on RNM signals with no threshold i.e. 0%, as in Tables VII, VIII, IX, are employed over 1 day. The Sharpe ratios of the trading results based on daily trading profits are reported as follows. Following Neumann et.al. we perform similar bootstrap and report the 95% confidence in the pair of numbers in the brackets below the averaged Sharpe ratio. Our results show no profitability net of trading cost at 1% significance level whereas Neumann et.al. (Table 10) displayed for comparison shows profitability before deducting transaction costs. They are not significantly larger than 0 at 1-tail 2.5% level, but are reported in Nuemann et.al. as significantly larger than 0, before spread costs, at 1-tail 5% level. The conclusions from both studies with respect to daily trading profitability look similar.

Skewness Trading Strategies			
Our Forecasting Models	Sharpe Ratio	N&S Forecasting Models	Sharpe Ratio
ARMA	0.94 (-0.62,2.44)	ARIMA(1,1,1)	2.58 (2.03,3.15)
AR(G)	0.70 (-0.85,2.18)	ARIMAX(1,1,1)	2.81 (2.24,3.39)
VECM	0.24 (-1.36,1.72)	VECM(1)	1.94 (1.35,2.50)
VAR(G)	0.83 (-0.72,2.33)	VECM-X(1)	2.16 (1.59,2.77)
Kurtosis Trading Strategies			
Our Forecasting Models	Sharpe Ratio	N&S Forecasting Models	Sharpe Ratio
ARMA	2.30 (-2.26,7.23)	ARIMA(1,1,1)	-0.22 (-0.64,0.32)
AR(G)	3.61 (-0.85,8.26)	ARIMAX(1,1,1)	-0.05 (-0.52,0.62)
VECM	-1.67 (-5.95,3.17)	VECM(1)	0.01 (-0.52,0.69)
VAR(G)	0.60 (-4.17,5.10)	VECM-X(1)	0.52 (-0.14,1.31)

Table XI

Intraday Ex-Ante Risk-Neutral Moments and Subsequent Returns

In single period asset pricing or intertemporal asset pricing with stationary stochastic investment opportunities, there are a majority of models where higher systematic variance or covariance, systematic co-kurtosis, and lower co-skewness are compensated by higher expected returns. These measures pertain to every security under equilibrium pricing, and are measures related to variance, skewness and kurtosis, except the latter can be diversified in a portfolio but the former cannot. In a somewhat similar setup, portfolios of stocks that are left-skewed, having high variance, and high kurtosis, can achieve a certain degree of within portfolio diversification (though not the general equilibrium model diversification), and thus across these portfolios, their average returns will display higher ex-post return compensation. The foregoing relationships is about equilibrium asset pricing and does not directly have anything to do with informational market efficiency. However, they are consistent with the idea that for each security, the contemporaneous return is related to its other moments; if idiosyncratic noise for a stock is small, then clearly this linear relationship will allow for the moments to be significantly priced (as in factor pricing models). Such a linear representation of return on its own moments is of course an approximate proxy for an equilibrium model. Moreover, if moments (we consider RNMs) can be forecasted as in Table II or for the more sophisticated ones in Tables IV to VI, then we can write: $\ln\left(\frac{F_{t+1}(\tau)}{F_t(\tau)}\right) = \beta_0 + \beta_1 \ln\left(\frac{F_t(\tau)}{F_{t-1}(\tau)}\right) + \beta_2 \Delta RN M_t(\tau) + \epsilon_{t+1}(\tau)$, where $t = 1, \dots, T$ and T is the number of 10-minutes interval in any given day. $\epsilon_{t+1}(\tau)$ is the i.i.d. residual error. τ indicates the correspondence of the price and RNM to futures and futures options with maturity τ . Results are displayed for $\tau = 1, 5$ days. (Results for other τ 's are basically similar.) For each trading day where there are many trades throughout the day, the regression is performed and the coefficient estimates are collected. Over the total number of such trading days, the estimated coefficients are averaged to obtain their means. Their t-statistics, within brackets reflecting if their average deviates significantly from zero, are also evaluated and reported in this table. The results for the above regressions are reported in each column for different models 1 to 4. Models 1,2, and 3 use RNV,RNS, and RNK respectively. Model 4 uses all the RNMs as explanatory variables. Bootstrap is carried out over 2000 iterations of the sample size indicated by the total number of trades. **, *, + denote 2-tail 5%, 2%, 1% significance levels respectively.

τ	Coefficient Estimate (Mean)	Model 1	Model 2	Model 3	Model 4
1 day	$\hat{\beta}_0$	0.00002	0.00002	0.00002	0.00000
	(t-stats)	(0.3172)	(0.2429)	(0.2780)	(0.0747)
	$\hat{\beta}_1$	-0.1110 ⁺	-0.0944 ⁺	-0.0922 ⁺	-0.1183 ⁺
	(t-stats)	(-4.0207)	(-3.1075)	(-3.1210)	(-3.9444)
	$\hat{\beta}_2$	-0.1020	-0.0001	0.0000	-0.2453
	(t-stats)	(-1.1207)	(-0.6488)	(0.7261)	(-1.6663)
	$\hat{\beta}_3$				-0.0003
	(t-stats)				(-1.5747)
5 days	$\hat{\beta}_4$				-0.0000
	(t-stats)				(-0.3899)
	$\hat{\beta}_0$	-0.00001	-0.00001	-0.00000	-0.00002
	(t-stats)	(-0.2228)	(-0.1324)	(-0.0144)	(-0.3485)
	$\hat{\beta}_1$	-0.1297 ⁺	-0.1335 ⁺	-0.1287 ⁺	-0.1381 ⁺
	(t-stats)	(-3.0784)	(-2.8033)	(-3.007)	(-2.5077)
	$\hat{\beta}_2$	-0.0211	0.0001	0.0000	-0.0086
	(t-stats)	(-0.6609)	(0.3918)	(0.1795)	(-0.2253)
$\hat{\beta}_3$				0.0003	
(t-stats)				(0.6547)	
	$\hat{\beta}_4$				0.0001
	(t-stats)				(0.5645)

References

- Aït-Sahalia, Y., Wang, Y., Yared, F., 2001. Do option markets correctly price the probabilities of movement of the underlying asset? *Journal of Econometrics* 102, 67–110.
- Amaya, D., Christoffersen, P., Jacobs, K., Vasquez, A., 2015. Does realized skewness predict the cross-section of equity returns? *Journal of Financial Economics* 118, 135–167.
- Andersen, T. G., Bollerslev, T., Diebold, F. X., Labys, P., 2003. Modeling and forecasting realized volatility. *Econometrica* 71, 579–625.
- Bakshi, G., Kapadia, N., Madan, D., 2003. Stock return characteristics, skew laws and the differential pricing of individual equity options. *Review of Financial Studies* 16, 101–143.
- Bali, T. G., Hu, J., Murray, S., 2013. Option implied volatility, skewness, and kurtosis and the cross-section of expected stock returns. Working Paper .
- Bali, T. G., Murray, S., 2013. Does risk-neutral skewness predict the cross section of equity option portfolio returns? *Journal of Financial and Quantitative Analysis* 48, 1145–1171.
- Barndorff-Nielsen, O., Nicolato, E., Shephard, N., 2002. Some recent developments in stochastic volatility modelling. *Quantitative Finance* 2, 11–23.
- Bollerslev, T., 1986. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* 31, 307–328.
- Breeden, D. T., Litzenberger, R. H., 1978. Prices of state-contingent claims implicit in option prices. *The Journal of Business* .
- Brooks, C., Burke, S. P., Heravi, S., Persaud, G., 2005. Autoregressive conditional kurtosis. *Journal of Financial Econometrics* 3, 399–421.
- Carr, P., Madan, D., 2001. Optimal positioning in derivative securities. *Quantitative Finance* 1, 19–37.
- Chang, B., Christoffersen, P., Jacobs, K., 2013. Market skewness risk and the cross section of stock returns. *Journal of Financial Economics* 107, 46–68.
- Chen, Y., Wolfgang, K., Pigorsch, U., 2010. Localized realized volatility modeling. *Journal of the American Statistical Association* 105, 1376–1393.

- Conrad, J., Dittmar, R. F., Ghysels, E., 2013. Ex ante skewness and expected stock returns. *The Journal of Finance* 68, 85–124.
- Cremers, M., Halling, M., Weinbaum, D., 2013. Aggregate jump and volatility risk in the cross-section of stock returns. Working Paper .
- Dennis, P., Mayhew, S., 2002. Risk-neutral skewness: Evidence from stock options. *Journal of Financial and Quantitative Analysis* 37, 471–493.
- Dittmar, R., 2002. Nonlinear pricing kernels, kurtosis preference, and evidence from the cross section of equity returns. *Journal of Finance* 57, 369–403.
- Efron, B., 1979. Bootstrap methods: Another look at the jackknife. *The Annals of Statistics* 7, 1–26.
- Engle, R. F., 1982. Autoregressive conditional heteroskedasticity with estimates of the variance of united kingdom inflation. *Econometrica* 50, 987–1007.
- Figlewski, S., 2008. Estimating the implied risk neutral density for the U.S. market portfolio. In: Bollerslev, T., Russell, J. R., Watson, M. (eds.), *Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle*, Oxford University Press.
- Figlewski, S., 2013. Some results on extracting and understanding the risk neutral returns distribution for the U.S. stock market. In: *NYU Stern–Federal Reserve Conference on Risk Neutral Densities*.
- Friesen, G. C., Zhang, Y., Zorn, T. S., 2012. Heterogeneous beliefs and risk-neutral skewness. *Journal of Financial and Quantitative Analysis* 47, 851–872.
- Gârleanu, N., Pedersen, L. H., Poteshman, A. M., 2009. Demand-based option pricing. *The Review of Financial Studies* 22, 4259–4299.
- Ghysels, E., Harvey, A., Renault, E., 1996. Stochastic volatility. In: Maddala, G., Rao, C. (eds.), *Statistical Methods in Finance*, North Holland Amsterdam, pp. 119–191.
- Gonçalves, S., Guidolin, M., 2006. Predictable dynamics in the s&p 500 index options implied volatility surface. *Journal of Business* 79, 1591–1635.
- Harvey, C., Siddique, A., 1999. Autoregressive conditional skewness. *Journal of Financial and Quantitative Analysis* 34, 465–487.
- Harvey, C., Siddique, A., 2000. Conditional skewness in asset pricing tests. *Journal of Finance* 55, 1263–1296.

- Jackwerth, J., Rubinstein, M., 1996. Recovering probability distributions from option prices. *Journal of Finance* 51, 1611–1631.
- Jiang, G. J., Tian, Y. S., 2005. The model-free implied volatility and its information content. *Review of Financial Studies* 18, 1305–1342.
- Kraus, A., Litzenberger, R., 1976. Skewness preference and the valuation of risk assets. *Journal of Finance* 31, 1085–1100.
- Merton, R. C., 1973. Theory of rational option pricing. *Journal of Economics and Management Science* 4(1), 141–183.
- Neumann, M., Skiadopoulos, G., 2013. Predictable dynamics in higher-order risk-neutral moments: Evidence from the s&p 500 options. *Journal of Financial and Quantitative Analysis* 48, 947–977.
- Poon, S.-H., Granger, C. W., 2003. Forecasting volatility in financial markets: A review. *Journal of Economic Literature* 41, 478–539.
- Ratcliff, R., 2013. Relative option prices and risk-neutral skew as predictors of index returns. *The Journal of Derivatives* pp. 89–105.
- Rubinstein, M., 1996. Implied binomial trees. *Journal of Finance* 49, 771–818.
- Taylor, S., Yadav, P. K., Zhang, Y., 2009. Cross-sectional analysis of risk-neutral skewness. *Journal of Derivatives* 16, 38–52.