Journal of Financial Economics 108 (2013) 409–424

Contents lists available at SciVerse ScienceDirect

Journal of Financial Economics

journal homepage: www.elsevier.com/locate/jfec

Risk and return: Long-run relations, fractional cointegration, and return predictability ☆

Tim Bollerslev [a,b,c,*], Daniela Osterrieder [c], Natalia Sizova [d], George Tauchen [a]

[a] Department of Economics, Duke University, Durham, NC 27708, USA
[b] NBER, Cambridge, MA 02138, USA
[c] CREATES, Department of Economics and Business, Aarhus University, 8210 Aarhus V, Denmark
[d] Department of Economics, Rice University, Houston, TX 77251, USA

ARTICLE INFO

Article history:
Received 30 December 2011
Received in revised form
12 July 2012
Accepted 16 July 2012
Available online 9 January 2013

JEL classification:
C22
C32
C51
C52
G12
G14

Keywords:
High-frequency data
Realized and options implied volatilities
Volatility risk premium
Long-memory and fractional cointegration
Return predictability

ABSTRACT

Univariate dependencies in market volatility, both objective and risk neutral, are best described by long-memory fractionally integrated processes. Meanwhile, the ex post difference, or the variance swap payoff reflecting the reward for bearing volatility risk, displays far less persistent dynamics. Using intraday data for the Standard & Poor's 500 and the volatility index (VIX), coupled with frequency domain methods, we separate the series into various components. We find that the coherence between volatility and the volatility-risk reward is the strongest at long-run frequencies. Our results are consistent with generalized long-run risk models and help explain why classical efforts of establishing a naïve return-volatility relation fail. We also estimate a fractionally cointegrated vector autoregression (CFVAR). The model-implied long-run equilibrium relation between the two variance variables results in nontrivial return predictability over interdaily and monthly horizons, supporting the idea that the cointegrating relation between the two variance measures proxies for the economic uncertainty rewarded by the market.

1. Introduction

We develop a unified framework for jointly modeling the dynamic dependencies and interrelatedness in aggregate stock market returns, realized volatilities, and options implied volatilities. Our estimation results rely on newly available high-frequency intraday data for the Standard & Poor's (S&P) 500 market portfolio and the corresponding Chicago Board Options Exchange (CBOE) Market Volatility Index (VIX), along with frequency domain inference procedures that allow us to focus on specific dependencies in the data. Our formal model setup is based on the co-fractional vector autoregression (VAR) of Johansen (2008a,b). We show that the longer-run dependencies inherent in the high-frequency data are consistent with the implications

☆ We would like to thank our referee, Neil Shephard, for very helpful comments that greatly improved the paper. We would also like to thank seminar participants at the 2011 Stanford Institute for Theoretical Economics (SITE) Conference, the 2011 Center for Research in Econometric Analysis of Time Series (CREATES) Long-Memory Conference at the University of Aarhus, Denmark, and the 2011 Escola de Séries Temporais e Econometria (ESTE) conference in Gramado, Brazil, for helpful comments. Tim Bollerslev's work was supported by a grant from the National Science Foundation (NSF) to the National Bureau of Economic Research (NBER) and CREATES funded by the Danish National Research Foundation.

* Corresponding author at: Department of Economics, Duke University, Durham, NC 27708, USA.
E-mail address: boller@duke.edu (T. Bollerslev).

0304-405X/$ - see front matter
http://dx.doi.org/10.1016/j.jfineco.2013.01.002

from the stylized equilibrium model in Bollerslev, Sizova and Tauchen (2012) that directly links the dynamics of the two volatility measures and the returns. Further corroborating the qualitative implications from that same theoretical model, we show that the variance risk premium estimated as the long-run equilibrium relation within the fractionally cointegrated system results in nontrivial return predictability over longer interdaily and monthly return horizons.

An enormous empirical literature has been devoted to characterizing the dynamic dependencies in stock market volatility and the linkages between volatilities and returns. The most striking empirical regularities to emerge from this burgeon literature are that volatility appears to be highly persistent, with the longer-run dependencies well described by a fractionally integrated process (see, e.g., Ding, Granger, and Engle, 1993; Baillie, Bollerslev, and Mikkelsen, 1996; Andersen and Bollerslev, 1997a; Comte and Renault, 1998); volatilities implied from options prices typically exceed the corresponding subsequent realized volatilities, suggesting that the reward for bearing pure volatility risk is negative on average (see, e.g., Bakshi and Kapadia, 2003; Carr and Wu, 2009; Bollerslev, Gibson, and Zhou, 2011); the volatility risk premium, defined as the difference between options implied and realized volatilities, tends to be much less persistent than the two individual volatility series, pointing to the existence of a fractional cointegration-type relation (see, e.g., Christensen and Nielsen, 2006; Bandi and Perron, 2006); volatility responds asymmetrically to lagged negative and positive returns, typically referred to as a leverage effect (see, e.g., Black, 1976; Nelson, 1991; Bollerslev, Litvinova, and Tauchen, 2006); and counter to the implications from a traditional risk-return trade-off-, or volatility feedback-, type relation, returns are at best weakly positively related, and sometimes even negatively related, to past volatilities (see, e.g., French, Schwert, and Stambaugh, 1987; Glosten, Jagannathan, and Runkle, 1993; Campbell and Hentschell, 1992).

The co-fractional VAR for the S&P 500 returns, realized volatilities, and VIX developed here is generally consistent with all of these empirical regularities. In contrast to most of the studies cited above, which are based on daily or coarser sampled data, our use of high-frequency five-minute observations on returns *and* volatilities allows for much sharper empirical inference concerning the second-order dynamic dependencies and frequency-specific linkages between the different variables. This, in turn, helps us to identify the periodicities in volatility and risk premia that are likely more important economically.

Our formal theoretical motivation for the use of high-frequency data are essentially twofold. First, the spectral density on $[-\pi,\pi]$ of a discretely sampled continuous time covariance stationary process could be seen as a folded-up version of the spectral density of the underlying continuous time process.[1] If the five minutes are thought of as very close to continuous time, the distortions invariably induced by the use of discretely sampled data are, therefore, minimal. In the same vein, the estimation in the time domain of the co-fractional VAR captures nearly all dynamics with arguably minimal distorting aliasing effects. Second, the coefficients of the co-fractional VAR are defined by second-order moments. These are known to be more precisely estimated the more finely sampled the data, so long as the number of parameters to be estimated remains small relative to the size of the data. As such, the efficiency of the inference is generally enhanced by the use of high-frequency data.

The plan for the rest of the paper is as follows. Section 2 provides a description of the high-frequency data underlying our empirical investigations. Section 3 characterizes the long-run dynamic dependencies in the two variance series, including the variance risk premium and the evidence for fractional cointegration. Section 4 details the risk-return relations inherent in the high-frequency data based on different variance proxies across different frequency bands. These results, in turn, motivate our empirical implementation of the fractionally cointegrated VAR system discussed in Section 5. Section 6 concludes.

2. Data

Our analysis is based on high-frequency tick-by-tick observations on the Chicago Mercantile Exchange (CME) futures contract for the S&P 500 aggregate market portfolio and the corresponding CBOE VIX volatility index. The data were obtained from Tick Data Inc. and cover the period from September 22, 2003 to January 31, 2012 and to December 30, 2011, respectively; further details concerning the data are available in Bollerslev, Sizova, and Tauchen (2012), where the same data have been analyzed from a different perspective over a shorter time period. Following standard practice in the literature as a way to guard against market-micro structure contaminants (see, e.g., Andersen, Bollerslev, Diebold, and Ebens, 2001), we transform the original tick-by-tick data into equally spaced five-minute observations using the last price within each five-minute interval. We denoted the corresponding geometric (log) returns by

$$r_{t+1} = \log(P_{t+1}) - \log(P_t), \tag{1}$$

where the time subscript t refers to the 77 intraday return observations plus the one overnight return per trading day. We define the corresponding risk-neutral return variation from the CBOE VIX as follows:

$$VIX_t^2 = \tfrac{30}{365}(VIX_t^{CBOE})^2, \tag{2}$$

where the scaling by 30/365 transforms the squared annualized observations on the CBOE index into monthly units.

It is well known that return volatility varies over the trading day and that it is typically higher around the open and close of trading than during the middle of the day (see, e.g., Andersen and Bollerslev, 1997b). The average variation of the overnight returns also far exceeds that of the intraday five-minute returns. This diurnal pattern in the return volatility could in turn result in a few observations each day dominating the analysis. To guard against this, we further standardize the raw returns in Eq. (1) by

[1] More technically, the spectral density of a discrete time scalar process could be thought of as identical copies over $(-\infty,\infty)$ on intervals $[(2k-1)\pi,(2k+1)\pi]$, $k=0,\pm 1,\pm 2,\ \ldots$, with the obvious analogy in the m-dimensional vector case on $[(2k-1)\pi,(2k+1)\pi]^m$.

Table 1
Summary statistics.

The table reports standard summary statistics for the returns, r_t, monthly realized variances, rv_t, and risk-neutral variances, vix_t^2. All of the statistics are based on five-minute observations from September 22, 2003 through December 30, 2011, for a total of 162,786 observations.

	Mean	Standard deviation	Autocorrelations					
			1	2	3	10	78	78×22
r_t	−0.0000	0.1478	−0.0343	−0.0151	−0.0075	−0.0041	−0.0016	0.0039
rv_t	2.9914	0.9948	1.0000	1.0000	1.0000	0.9999	0.9981	0.8045
vix_t^2	3.4133	0.8259	0.9999	0.9997	0.9996	0.9985	0.9884	0.8630

the sample standard deviations for each of the 78 time intervals. All of the estimation results reported on below are based on these standardized returns. However, when assessing the return predictability implied by the different models, we always report the relevant measures in terms of the unstandardized returns of practical interest.

Following the simple realized volatility estimators of Andersen, Bollerslev, Diebold, and Labys (2001) and Barndorff-Nielsen and Shephard (2002), a number of alternative nonparametric procedures have been developed for more efficiently estimating the ex post return variation from high-frequency data, including the realized kernels of Barndorff-Nielsen, Hansen, Lunde, and Shephard (2008), the multi-scale estimators of Zhang, Mykland, and Aït-Sahalia (2005), and the pre-averaging estimator of Jacod, Li, Mykland, Podolskij, and Vetter (2009). Here we rely on a simple-to-implement subsampling and averaging estimator for the one-month (or 22 trading days) ex post return variation obtained by summing the within month five-minute and overnight squared returns. Specifically,

$$RV_t = \frac{1}{5}\sum_{j=1}^{5}\sum_{i=1}^{78\times22} r^2_{t+i+(j-1)/5}, \tag{3}$$

where $r_{t+i+(j-1)/5}$ for $j=1,\ldots,5$ refer to the five different five-minute return series obtained by initializing the day at five different one-minute time stamps. This averaging of the five separate squared return series naturally enhances the efficiency of the estimator relative to the basic realized volatility estimator based on a single five-minute return series, or $j=1$. The one-month forward horizon of RV_t, as defined in Eq. (3), exactly matches the return variation prized by the VIX_t^2, as defined in Eq. (2). As such, the difference between VIX_t^2 and RV_t corresponds directly to the ex post payoff from selling a variance swap contract.

Finally, to help stabilize the two variance measures and render them more amenable to linear time series modeling, we transform both into logarithmic units:

$$rv_t = \log(RV_t), \tag{4}$$

$$vix_t^2 = \log(VIX_t^2). \tag{5}$$

With the loss of one month at the end of the sample due to the calculation of rv_t, our common data set for the three series thus covers the period from September 22, 2003 to December 30, 2011, for a total of $T=2087\times78=162{,}786$ observations; comparable results based on a shorter sample through December 31, 2008 are available in an earlier working paper version of the paper.

Standard summary statistics for each of the three series are reported in Table 1. The high-frequency returns are approximately serially uncorrelated, with a mean indistinguishable from zero. Consist with the extant literature, the unconditional mean of the realized variance is lower than the mean of the risk-neutral variance, indicative of an on average positive risk premium for bearing volatility risk. At the same time, the risk-neutral variance appears slightly less volatile than the realized variance. Both of the variance series exhibit substantial persistence with extremely slow decay in their autocorrelations. The next section further details these dynamic dependencies in the two variance series.

3. Variance dynamics

3.1. Long-run volatility dependencies

The notion of fractional integration often provides a convenient statistical framework for capturing long-run dependencies in economic time series (see, e.g., the discussion in Baillie, 1996). A stationary time series y_t is said to be fractionally integrated of order $d\in(0,0.5)$, written $I(d)$, if

$$\Delta^d y_t = e_t, \tag{6}$$

where e_t is an $I(0)$ process and $\Delta^d=(1-L)^d$ denotes the fractional difference operator

$$\Delta^d \equiv (1-L)^d = \sum_{i=0}^{\infty}(-1)^i\binom{d}{i}L^i. \tag{7}$$

The spectral density of the process y_t has a pole of the order ω^{-2d} for frequency ω near the origin, while the filtered series $\Delta^d y_t$ has finite spectral density at the origin.

Ample empirical evidence shows that financial market volatilities are well described by covariance stationary $I(d)$ processes with fractional integration parameter close to but less than one-half.[2] For instance, Andersen, Bollerslev, Diebold, and Ebens (2001) report average

[2] One notable exception is Bandi and Perron (2006), who argue that d is in excess of 0.5 and, thus, outside the stationary region. There is also a long history of using superpositions of short-memory processes as a way to capture long-range dependence in realized volatilities (see, e.g., Barndorff-Nielsen and Shephard, 2002; Corsi, 2009).

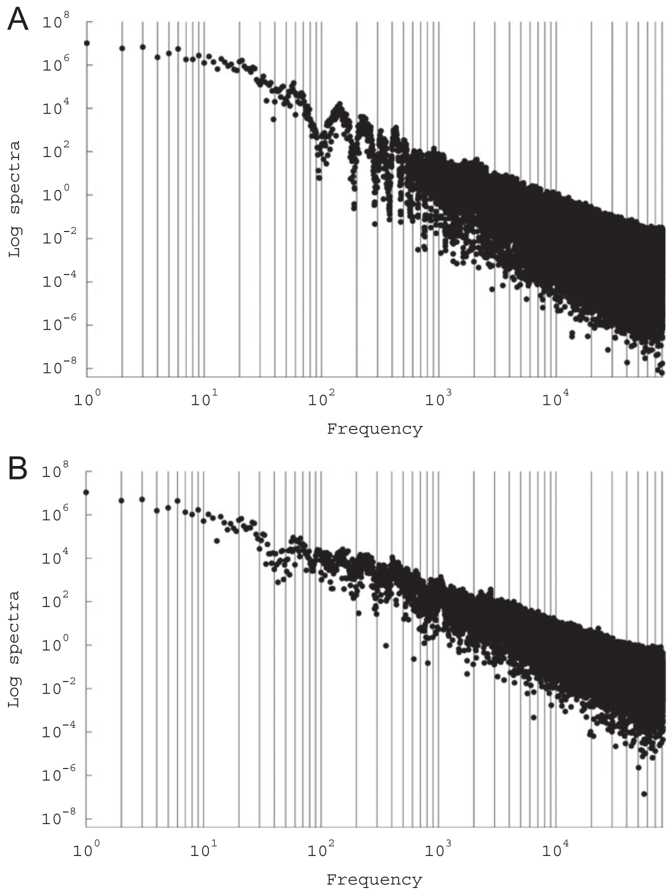

Fig. 1. Periodograms for realized and risk-neutral variances. The figure plots the sample periodograms of the realized variance rv_t (Panel A) and the risk-neutral variance vix_t^2 (Panel B). The periodograms are plotted on a double logarithmic scale. The estimates are based on five-minute observations from September 22, 2003 to December 30, 2011.

fractional integration parameters for a set of realized equity return volatilities of approximately 0.35, and the results for foreign exchange rates in Andersen, Bollerslev, Diebold, and Labys (2003) suggest that d is close to 0.4. Similarly, Christensen and Nielsen (2006) find that daily realized and options implied equity index volatilities are fractionally integrated with d around 0.4.

Most of the existing literature, including the above-cited studies, have relied on daily or coarser sampled realized and options implied volatility measures for determining the order of fractional integration.[3] By contrast, both of our volatility series are recorded at a five-minute sampling frequency. Hence, as a precursor to our more detailed joint empirical analysis, we begin by double-checking that the folding of the spectral densities associated with the lower sampling frequency have not distorted the previously reported estimates for d.

To this end, Fig. 1 shows the raw log-periodograms of the five-minute vix_t^2 and rv_t series at the harmonic frequencies $\omega_j = (2\pi/T)j$, $j = 1, 2, \ldots, T/2$. The periodograms of the two variance variables are similar, with most of the power concentrated at the low frequencies. At the same time, there appears to be three distinct regions within the frequency domain: a relatively short left-most low-frequency region up until $\underline{\omega} \approx 0.0010$, where the log-periodograms are linear and nearly flat; an intermediate region between $\underline{\omega}$ and $\overline{\omega} \approx 0.0806$, with more steeply sloped periodograms; and a third right-most region to the right of $\overline{\omega}$ corresponding to the within-day variation in the volatilities, where the periodograms are erratic.

The approximate linearity of the log-periodograms for the low frequencies close to zero directly points to long-memory dependencies, or fractional integration. We estimate the fractional integration parameter d using both the log-periodogram regression of Geweke and Porter-Hudak (1983) and the local-Whittle likelihood procedure of Künsch (1987). For both estimators, we set the required truncation parameter to $j_{max} = 26$, corresponding to $\underline{\omega}$ and periodicities of 3.5 months and longer being used in the estimation. Based on visual inspection of Fig. 1 this seemingly covers the frequency range where the long-memory behavior hold true; see also the related discussion in Sowell (1992). The resulting estimates for d, with

[3] Some studies have also sought to estimate d based on high-frequency squared, or absolute, returns interpreted as noisy proxies for the true volatility. See, e.g., Andersen and Bollerslev (1997a) and Ohanissian, Russell, and Tsay (2009) and the many references therein.

asymptotic standard errors in parentheses, for the vix_t^2 series are 0.398 (0.126) using the log-periodogram and 0.416 (0.098) using the local Whittle approach. For the rv_t series, the same two estimates are 0.399 (0.126) and 0.403 (0.098), respectively. As such, our findings for the two five-minute volatility series are entirely consistent with the typical estimates for d reported in the existing literature.

3.2. Fractional cointegration in variances

The preceding results strongly suggest that each of the two high-frequency variance series are individually long-memory processes. At the same time, one might naturally expect that the two variance series are tied together in the long run in the form of fractionally cointegrated-type relation. Following Granger (1986), if a linear combination of two fractionally integrated variables is integrated of a lower order than those of the individual series, then the variables are said to be fractionally cointegrated; see also Robinson and Marinucci (2003).

The simplest case of fractional cointegration occurs when the two individual series share the same order of fractional integration, but their difference is integrated of a lower order. In the present context, this difference naturally corresponds to the ex post payoff from selling a variance swap, $VIX_t^2 - RV_t$, or in logarithmic form

$$vd_t = vix_t^2 - rv_t. \tag{8}$$

The fact that vd_t exhibits less persistence than the two individual variance series has previously been documented with daily and lower frequency data by Christensen and Nielsen (2006), Bandi and Perron (2006), and Chernov (2007), among others.[4]

To establish a similar long-run relation between our two high-frequency variance series, we begin by testing for equality of the fractional difference parameters using the Wald test of Robinson (1995). The asymptotically distributed χ_1^2 test statistic equals 0.018 when the ds are estimated by the local Whittle procedure and 0.000 for the d's estimated by the log-periodogram estimator. Either way, little evidence seems to exist against the hypothesis that the two variance series are fractionally integrated of the same order.

To further explore the possibility of fractional cointegration, we next consider the following linear regression:

$$rv_t = \beta_0 + \beta_1 vix_t^2 + \nu_t. \tag{9}$$

This regression is analogous to the Mincer-Zarnowitz style regression commonly used for evaluating the quality of macroeconomic time series forecasts. That is, it evaluates whether vix_t^2 is conditionally unbiased for the ex post realized variance rv_t. The residuals from this regression reduce to the variance difference defined in Eq. (8) for $\beta_0 = 0$ and $\beta_1 = 1$. However, instead of restricting the relation between rv_t and vix_t to be the same across all frequencies, we estimate the regression using low-pass frequency domain least squares (FDLS) (see, e.g., Robinson, 1994). As before, we truncate the regression at $j_{max} = 26$, corresponding to periodicities of 3.5 months and longer. We rely on the local Whittle approach for estimating d for the error term ν_t, together with the techniques developed by Christensen and Nielsen (2006) and Shimotsu and Phillips (2006) for estimating the asymptotic standard errors.

The resulting FDLS estimate for the degree of fractional integration of the residuals equals $\hat{d}^{(\nu)} = 0.093(0.110)$, clearly supporting the idea that the high-frequency vix_t^2 and rv_t series are fractionally cointegrated. Moreover, the FDLS estimate of β_1 equals 1.090(0.116), slightly larger than unity but insignificantly so.

3.3. Dynamic dependencies across frequencies

In addition to assessing the integration order of rv_t and vix_t^2, and the long-run relation between the two variance measures, the joint distribution of the variance measures could be illuminated by decomposing each of the variables into their long-run, intermediate, and short-run components. In order to do so, we rely on time-domain band-pass filters for extracting the specific periodicities from the observed series.[5]

Fig. 1 reveals a change in the slopes of the periodograms for the realized and risk-neutral variances for frequencies around $\underline{\omega} \approx 0.0010$, or periods around 3.5 months. Motivated by this observation, we, therefore, extract a low-frequency, or long-run, component corresponding to dependencies in excess of 3.5 months. That is, for frequency $\underline{\omega}$ we define a lowpass-filtered series by

$$y_t^{(low)} = \sum_{i=-k}^{k} a_i L^i y_t, \tag{10}$$

where

$$a_i = \begin{cases} \dfrac{\sin(i\underline{\omega})}{i\pi} - \left(\dfrac{\underline{\omega}}{\pi} + 2\displaystyle\sum_{j=1}^{k} \dfrac{\sin(j\underline{\omega})}{j\pi} - 1 \right) \Big/ (2k+1), & i = \pm 1, \ldots, \pm k, \\ 1 - \displaystyle\sum_{h=-k}^{-1} a_h - \sum_{h=1}^{k} a_h, & i = 0. \end{cases} \tag{11}$$

Intuitively, the $y_t^{(low)}$ series essentially equates the part of the spectrum in Fig. 1 to the right of $\underline{\omega}$ to zero.

Looking at the two sample periodograms in Fig. 1, a distinct change is evident in the general patterns for frequencies in excess of the daily frequency $\overline{\omega} \approx 0.0806$.[6] To isolate the interdaily component with

[4] The overlapping nature of rv_t invariably generates substantial serial correlation in the high-frequency vd_t series.

[5] Compared with frequency-domain filters (see, e.g., Hassler, Lundvik, Persson, and Söderlind, 1994), the time-domain filters applied here have the advantage that the filtered series are time invariant and do not depend on the length of the sample. Similar filters to the ones used here have previously been applied by Andersen and Bollerslev (1997a) for decomposing high-frequency foreign exchange rates into interday and intradaily components. For additional discussion of band-pass filtering see Baxter and King (1999), where the same techniques have been used for extracting business cycle components of macroeconomic times series.

[6] It is well known that the volatilities of high-frequency returns, or point-in-time volatilities, exhibit strong U-shaped patterns across the

periodicities of less than 3.5 months from the intraday dynamics, we, therefore, compute an intermediate-frequency series. Specifically, for frequencies in the band $\underline{\omega} < \omega < \overline{\omega}$, we define the band-pass filtered series by

$$y_t^{(band)} = \sum_{i=-k}^{k} (b_i - a_i) L^i y_t, \tag{12}$$

where

$$b_i = \begin{cases} \dfrac{\sin(i\overline{\omega})}{i\pi} - \left(\dfrac{\overline{\omega}}{\pi} + 2\sum_{j=1}^{k} \dfrac{\sin(j\overline{\omega})}{j\pi} - 1 \right) \Big/ (2k+1), & i = \pm 1, \ldots, \pm k, \\ 1 - \sum_{h=-k}^{-1} b_h - \sum_{h=1}^{k} b_h, & i = 0. \end{cases} \tag{13}$$

Finally, the high-pass filtered series corresponding to periodicities of a day and shorter is simply computed by

$$y_t^{(high)} = y_t - \sum_{i=-k}^{k} b_i L^i y_t. \tag{14}$$

By definition, $y_t \equiv y_t^{(low)} + y_t^{(pass)} + y_t^{(high)}$.

The higher the value of the truncation parameter k, the closer the gains of the approximate filters in Eqs. (10), (12), and (14) are to the ideal gains of zero for the frequencies that are filtered out and unity for the desired frequency bands. In the implementation reported on below, we set $k = 78 \times 22$ for all of the three filters, resulting in a loss of one month of observations at the beginning and the end of the sample.

We first focus on the dynamics of the long-run components. Assuming that the effects of market frictions and short-run fluctuations disappear in the long run, the dynamics of the low-pass filtered series should, therefore, more clearly reveal the underlying equilibrium relations between the variables. Fig. 2 plots the two filtered variance measures $rv_t^{(low)}$ and $vix_t^{2(low)}$ defined by applying the filter in Eq. (10). Consistent with the results in Section 3.2, the figure reveals strong co-movements between the two low-frequency variance components. Indeed, the sample correlation between the two low-pass filtered series equals $Corr(rv_t^{(low)}, vix_t^{2(low)}) = 0.916$.

We further detail the relation between vix_t^2 and rv_t across all frequencies through measures of their inter-relatedness, or coherence. Our estimates of the coherence measures are based on the classic Tukey-Hanning method; see, e.g., the discussion in Granger and Hatanaka (1964). The coherence is analogous to the square of the correlation between two series, taking values from zero (no relation) to one (perfect correlation). In contrast to the standard correlation coefficient, however, the coherence is a function of frequency. The coherence being close to zero for a certain frequency range thus indicates the absence of any relation between the two series across those periodicities. As follows from the first panel in Fig. 3 labeled "Total," the coherence between vix_t^2 and rv_t is close to one for the lowest frequencies but decreases to close to zero for the higher frequencies. The almost perfect dynamic relation between the vix_t^2 and rv_t series, therefore, holds only in the long run.

To get a more nuanced picture of these dependencies, the remaining panels in Fig. 3 show the coherence for the high-pass, band-pass, and low-pass filtered variance series. The coherence for the long-run components is large and generally in excess of 0.6, across all frequencies. The coherence for the band-pass filter variances is substantially lower and around 0.05, and the coherence for the high-frequency components is practically zero.

To interpret these patterns, it is instructive to think about vix_t^2 as the sum of the ex post realized volatility rv_t, a premium for bearing volatility risk, vrp_t, along with the corresponding forecast error, say ξ_t:

$$vix_t^2 = rv_t + vrp_t + \xi_t. \tag{15}$$

For low frequencies the coherence between vix_t^2 and rv_t is close to one, implying that the influences of ξ_t and vrp_t are both fairly minor. This is also consistent with the findings that rv_t is integrated of a higher order than $vd_t = \xi_t + vrp_t$ and that most of the low-frequency dynamics in vix_t^2 are due to changes in the volatility. This close coherence is broken at the intermediate and ultra-high intraday frequencies, in which most of the changes in vix_t^2 stem from changes in the risk premium or expectational errors vis-à-vis the future realized volatility. These same arguments also facilitate the interpretation of the empirical risk-return relations, which we discuss next.

4. Risk-return relations

A large empirical literature has been devoted to the estimation of risk-return trade-off relations in aggregate equity market returns (see, e.g., the discussion in Rossi and Timmermann, 2012, and the many references therein). Much of this research is motivated by simple dynamic capital asset pricing model (CAPM)-type reasoning along the lines of

$$E_t(r_{t+1}) = \gamma \sigma_t^2, \tag{16}$$

where σ_t^2 represents the local return variance and γ is interpreted as a risk-aversion parameter.[7] The actual estimation of this equation necessitates a proxy for σ_t^2. By far the most commonly employed empirical approach relies on the (G)ARCH-M model (Engle, Lilien, and Robins, 1987) for jointly estimating the conditional mean of the returns together with the conditional variance of the returns in place of $\sigma_t^2 = Var_t(r_{t+1})$. Instead, by relying on the variance measures analyzed above as directly observable proxies for risk, Eq. (16) can be estimated directly. The following subsection explores this idea in our high-frequency

(*footnote continued*)
trading day. See, e.g., Harris (1986) and Andersen and Bollerslev (1997b). By contrast, the two volatility series analyzed here both measure the variation over a month, and as such even though they behave differently intraday they do not show the same strong, almost deterministic, intraday patterns.

[7] The Merton (1980) model sometimes used as additional justification for this relation formally involves the excess return on the market. The requisite Jensen's correction term for the logarithmic returns analyzed here simple adds $1/2$ to the value of γ. Also, the risk-free rate is essentially zero at the five-minute level.

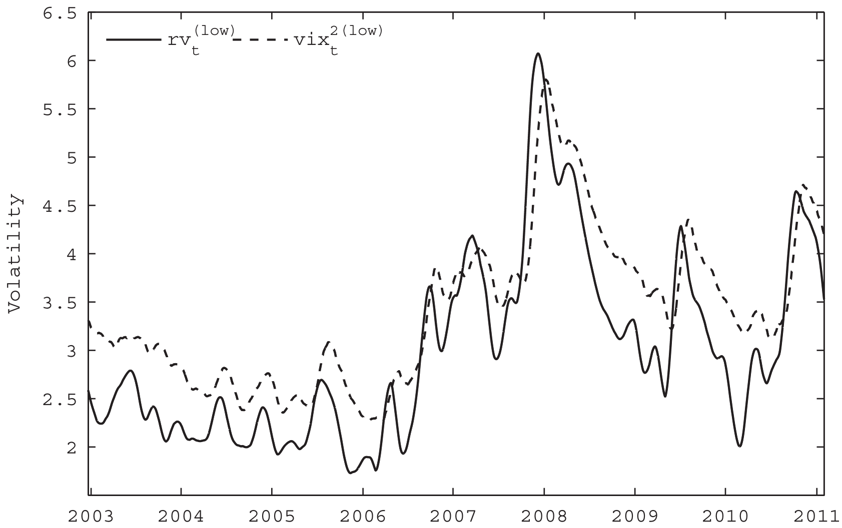

Fig. 2. Co-movement between low-frequency realized and risk-neutral variances. The figure plots the low-pass filtered realized variance rv_t (solid line) and risk-neutral variance vix_t^2 (dashed line) over the October 22, 2003 to November 28, 2011 sample period.

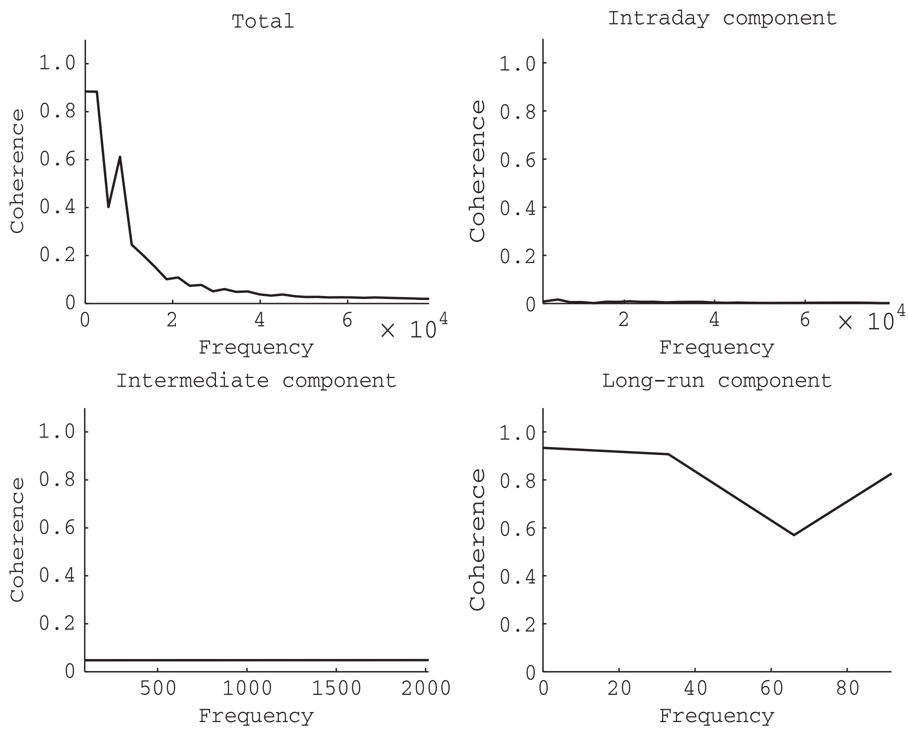

Fig. 3. Relation between realized and risk-neutral variances across frequencies. The figure plots the coherence between the realized variance rv_t and the risk-neutral variance vix_t^2, as well as their high-pass (less than one day), band-pass (one day to 3.5 months), and low-pass (longer than 3.5 months) filtered counterparts. All of the estimates are based on five-minute observations from October 22, 2003 to November 28, 2011.

data setting, keeping in mind the preceding characterization of the underlying volatility dynamics.[8]

[8] The use of high-frequency-based realized volatility measures in the estimation of a daily risk-return trade-off relation has previously been explored by Bali and Peng (2006).

4.1. Return-variance regressions

The basic risk-return relation in Eq. (16) can be conveniently expressed in regression form as

$$r_{t+1} = \alpha + \beta v_t + u_{t+1}, \tag{17}$$

where v_t denotes the specific variance proxy used in place of σ_t^2. The first two rows of Table 2 labeled "Raw" show

Table 2
Univariate return regressions.
The table reports five-minute return regressions, $r_t = \alpha + \beta v_{t-1} + u_t$, based on data from September 30, 2003 to December 30, 2011. The first two columns report the ordinary least squares estimates for β and the corresponding Newey-West standard errors (SE). The last column reports the regression R^2's. The fractional difference parameter is fixed at $d = 0.40$.

Risk proxy (v_t)	β	SE	R^2
Level			
rv_t	−0.00144	0.00066	0.0094%
vix_t^2	0.00028	0.00082	0.0002%
Long-memory adjusted			
$\Delta^d rv_t$	−0.0599	0.0123	0.0210%
$\Delta^d vix_t^2$	0.0480	0.0170	0.0151%
Variance difference			
$vd_t = vix_t^2 - rv_t$	0.00554	0.00077	0.0410%

the results of this regression fitted with our high-frequency data using either rv_t or vix_t^2 in place of v_t. The results are very weak. This is perhaps not surprising, in view of the difficulty to detect a significant relation in daily and coarser frequency data reported in the extant literature.

But these raw regressions are also unbalanced and likely not very informative. They involve an essentially white noise $I(0)$ variable on the left-hand side (the return) and a strongly persistent $I(d)$ variable on the right-hand side (the realized or risk-neutral variance). This same statistical problem has also previously been discussed in the context of regression-based tests for unbiasedness in the forward foreign exchange market by Baillie and Bollerslev (2000). Several means to cope with such unbalanced regressions have been suggested in the literature. Maynard, Smallwood, and Wohar (2013), in particular, focus explicitly on the case in which the predictor variable is fractionally integrated and propose an approach to rebalance the regression. A similar approach has also been applied by Christensen and Nielsen (2007), who consider a VAR framework in which the level of returns is predicted by the fractionally filtered variance series.

A simplified version of the return equation in Christensen and Nielsen (2007) is given by

$$r_{t+1} = \alpha + \beta \Delta^d v_t + u_{t+1}, \tag{18}$$

where $\Delta^d \equiv (1-L)^d$ denotes the fractional difference filter.[9] In implementing the fractional filters, we truncate the series expansion for $(1-L)^d v_t$ at the beginning of the sample, discarding the first week of filtered observations. Based on the results reported in Section 3 we fix $d \equiv 0.4$. The resulting long-memory-adjusted regressions using $\Delta^{0.4} rv_t$ or $\Delta^{0.4} vix_t^2$ as risk proxies are reported in the next two rows of Table 2. All of the regressions reported in the table are based on the identical sample period from September 30, 2003 to December 30, 2011.

[9] More precisely, Christensen and Nielsen (2007) consider a modified version of Eq. (18), where both r_{t+1} and v_t are detrended by their unconditional means and $\alpha \equiv 0$. In addition, their VAR structure allows past returns to affect current returns.

Compared with the "raw" regressions, both of the R^2's are somewhat higher. The estimated coefficient for the realized variance is now strongly statistically significant but signals a negative risk-return trade-off. By contrast, the fractionally differenced risk-neutral variance results in a statistically significant positive $\hat{\beta}$ estimate, albeit a lower R^2. These mixed results are perhaps again not surprising in view of the existing literature.

As highlighted in Section 3.2, another way to render the two variance series stationary, and, in turn, balance the regression, is to consider their difference. The last row in Table 2 labeled "Variance difference" reports the results from the regression considering vd_t as a proxy for risk,

$$r_{t+1} = \alpha + \beta vd_t + u_{t+1}. \tag{19}$$

The estimate of β is now positive and statistically significant. The regression R^2 is also larger than for any of the other regressions reported in the table. Thus, the difference between the two variance variables appears far more informative for the returns than each of the two variance variables in isolation, whether in their raw or fractionally filtered form. To help further gauge these simple regression-based results, it is informative to decompose the variables into their periodic components.

4.2. Risk-return relations across frequencies

Following the analysis in Section 3.3, the return r_t and the variance variables rv_t, vix_t^2, and vd_t are naturally decomposed via band-pass filtering into their short-, intermediate-, and long-run components. Table 3 summarizes the correlations between the resulting components for the

Table 3
Risk-return relations across frequencies.
The table reports the correlations between the short-, medium-, and long-run components of the returns r_t, the realized variance rv_t, the risk-neutral variance vix_t^2, and the variance difference vd_t. All of the correlations are based on five-minute observations from October 22, 2003 to November 28, 2011, along with the band-pass filtering procedures discussed in the text for decomposing the series into the different components. The standard errors (reported in parentheses) are based on heteroskedasticity and autocorrelation consistent (HAC) covariance matrix estimator, using a Bartlett kernel with bandwidth equal to $78 \times 5 = 390$.

	r_{t+1}^{high}	r_{t+1}^{band}	r_{t+1}^{low}
rv_t^{high}	−0.012 (0.003)		
rv_t^{band}		−0.074 (0.014)	
rv_t^{low}			−0.433 (0.064)
vix_t^{high}	0.173 (0.004)		
vix_t^{band}		−0.041 (0.020)	
vix_t^{low}			−0.193 (0.067)
vd_t^{high}	0.165 (0.004)		
vd_t^{band}		0.025 (0.016)	
vd_t^{low}			0.664 (0.035)

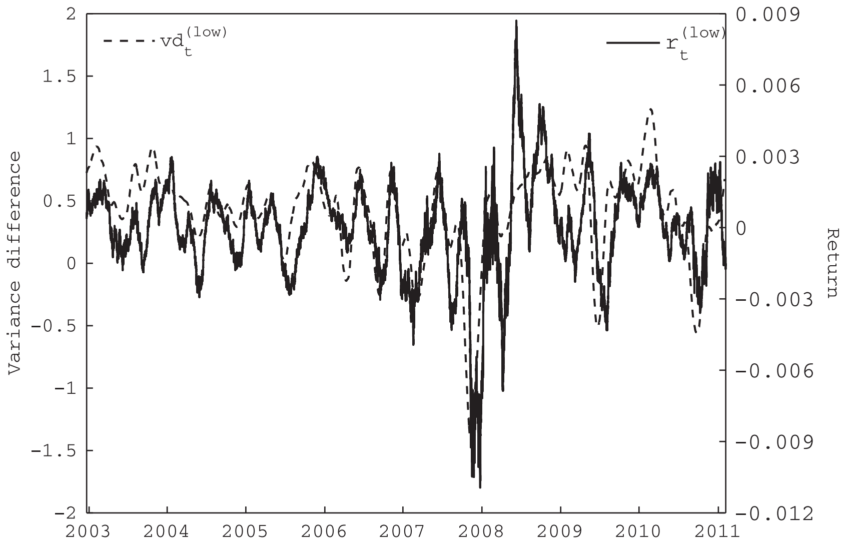

Fig. 4. Co-movement between low-frequency variance difference and returns. The figure plots the low-pass filtered variance difference vd_t (dashed line) and the returns r_t (solid line) for the October 22, 2003 to November 28, 2011 sample period.

different variance variables with the same components of the return. Evidently, not much of a relation exists at the high and intermediate frequencies. Again, this is not surprising as one might expect the association between risk and return to be more of a long-run than a short-run phenomenon. These results are also consistent with the monthly return regressions in Bollerslev, Tauchen, and Zhou (2009), which suggest that the return predictability of the variance risk premium is maximized at a four-month horizon. Four months lie within the frequency-band classified as low-frequency here.

Corroborating this conjecture, the low-frequency correlations reported in the right-most column in the table are much higher in magnitude for rv_t and vd_t and marginally higher for vix_t^2. Somewhat paradoxically, however, the correlations are negative for both of the individual variance variables, but positive and larger in magnitude for the variance difference. To visualize this relation, Fig. 4 shows the low-pass filtered $vd_t^{(low)}$ and $r_t^{(low)}$ series. The co-movement between the two series is apparent.

In a sum, neither of the two variance series, rv_t and vix_t^2, exhibit dynamic co-movements with the returns over any of the three frequency bands that would support the notion of a positive risk-return trade-off relation. If instead their difference is employed as a proxy for risk, the results are more in line with intuitive expectations regarding a positive low-frequency risk-return trade-off relation.

The continuous time model developed by Bollerslev, Sizova, and Tauchen (2012) is useful for interpreting these results. In the first place, the volatility σ_t^2 in the basic dynamic CAPM and Eq. (16) does not have the actual role of a volatility risk premium per se but can rather be viewed as a time-varying price on endowment (consumption) risk. The only variable in that simplistic framework that carries a risk premium is consumption, so, despite the intuitive appeal of model (16), there really is no reason to expect raw variance variables to relate to returns in any manner, except, perhaps, through some indirect mechanism that could be of either sign.[10] However, in a generalized long-run risk model with uncertainty about the variability of economic prospects (vol-of-vol), the difference in the variance variables, or the variance risk premium corresponding to vrp_t in Eq. (15), is the key factor connecting returns and variability. The difference is most highly associated with the vol-of-vol and overall economic uncertainty, and that factor commands a substantial risk premium by investors with recursive utility and a strong preference for early resolution of uncertainty. Consequently, the regression in Eq. (18) is entirely consistent with more sophisticated versions of the dynamic CAPM.

To fully explore these relations between the return and the two variance measures, and the implications thereof for return predictability, we next turn to the estimation of a joint model for r_t, rv_t, and vix_t^2, explicitly designed to accommodate the intricate dynamic and cross-variable dependencies highlighted so far.

5. Co-fractional system

The co-fractional VAR model of Johansen (2008a,b) affords a convenient statistical framework to distinguish long-run and short-run effects in a system setting involving fractionally integrated $I(d)$ variables. Specifically, let $z_t \equiv (rv_t, vix_t^2, r_t)'$ denote the five-minute 3×1 vector process. Guided by the empirical findings in the previous sections, the simplified versions of the co-fractional VAR model for z_t estimated here, say CFVAR$_d(p)$, takes the form

$$\Delta^d z_t = \gamma(\delta'(1-\Delta^d)z_t+\rho') + \sum_{i=1}^{p} \Gamma_i (1-\Delta^d)^i \Delta^d z_t + \epsilon_t, \qquad (20)$$

[10] The sign of the relation between returns and variances have also previously been called into question on theoretical grounds by Backus and Gregory (1993), among others.

where ϵ_t denotes a vector white noise process with unconditional covariance matrix Ω.[11]

This dynamic CFVAR representation directly parallels the classical error-correction-type representation with cointegrated $I(1)$ variables. The process z_t contains the fractionally integrated $I(d)$ variables, analogous to the $I(1)$-level variables in standard cointegration. The fractional difference operator $\Delta^d=(1-L)^d$ thus reduces the left-hand side of Eq. (20) from an $I(d)$ to an $I(0)$ process, just like the first-difference operator for standard cointegrated systems reduces the $I(1)$ variables to $I(0)$.[12] The right-hand side of the equation, therefore, must also be $I(0)$. The first term $(1-\Delta^d)z_t$ is what remains after applying the fractional difference operator and, thereby, must be $I(d)$. The matrix $\gamma\delta'$, therefore, has to be of reduced rank for this to be an $I(0)$ process. In this situation, $\delta'(1-\Delta^d)z_t$ has the interpretation of a (fractional) error-correction matrix, with γ the conformable matrix of impact coefficients. The second term on the right-hand side, involving the matrix fractional distributed lag and powers of $(1-\Delta^d)$ applied to $\Delta^d z_t$, directly mirrors the matrix distributed lag in standard error-correction models, in which powers of L are applied to first-differences of the underlying variables. The corresponding Γ_i matrices are essentially nuisance parameters, with p taken sufficiently large to render the disturbance term ϵ_t serially uncorrelated.

The empirical evidence presented in Section 3 suggests that rv_t and vix_t^2 are both $I(d)$ but that they fractionally cointegrate to an $I(0)$ process, while r_t is $I(0)$. Consequently, the column rank of $\gamma\delta'$ should be equal to two, with the natural normalization

$$\delta'=\begin{pmatrix}-\tilde{\delta} & 1 & 0\\ 0 & 0 & 1\end{pmatrix} \tag{21}$$

and the corresponding matrix of impact coefficients

$$\gamma=\begin{pmatrix}\gamma_{11} & \gamma_{12}\\ \gamma_{21} & \gamma_{22}\\ \gamma_{31} & \gamma_{32}\end{pmatrix} \tag{22}$$

left unrestricted. The parameter $\tilde{\delta}$ naturally governs the long-run tie between rv_t and vix_t^2, which in turn defines the variance error correction term as a linear combination of multiple lags of $(1-\Delta^d)(vix_t^2-\tilde{\delta}rv_t)$.[13]

The parameters γ_{11} and γ_{21} capture any internal long-run relations between the variance error-correction term and the variance variables themselves, while γ_{31} captures the relation between the variance error-correction term and the returns, or the long-run dynamic "volatility feedback" effect implied by the model. The corresponding short-run counterparts are determined by the $\Gamma_i^{(31)}$ and $\Gamma_i^{(32)}$ parameters. Similarly, the long-run dynamic "leverage effect" depends on the values of γ_{12} and γ_{22}, with the corresponding short-run effects determined by $\Gamma_i^{(13)}$ and $\Gamma_i^{(23)}$. The overall degree of return predictability implied by the model is jointly determined by the γ_{3j} and $\Gamma_i^{(3j)}$ parameters.

In sum, by estimating the CFVAR$_d(p)$ model we draw on a richer information set than we did in the previous sections. Importantly, by separately parameterizing the long-run and the short-run dynamics of returns and the variance series, the model is able to accommodate empirically realistic $I(d)$ long-memory in the realized and the risk-neutral variances and their fractional cointegration, while maintaining that the returns are $I(0)$.

5.1. Estimation results

To facilitate the estimation of the CFVAR$_d(p)$ model we begin by fixing the value of the fractional integration parameter d. The estimation then proceeds in two stages. In the first step, we use the preset fractional differencing parameter d to construct the vector time series $\Delta^d z_t$ of filtered realized variances, risk-neutral variances, and returns.[14] In the second stage, we obtain parameter estimates for the co-fractional model by iterated seemingly unrelated regression (SUR). We separately select the number of lags for each of the three equations, allowing for different number of lags of rv_t, vix_t^2, and r_t in each of the equations as determined by Schwarz's Bayesian Information Criterion (BIC). Finally, we select the value of d that maximizes the Gaussian likelihood function subject to the BIC chosen lag specifications.

Turning to the actual CFVAR estimation results reported in Table 4, we find that the likelihood function is maximized at $d=0.4$. This value of d is directly in line with the semiparametric estimates discussed in Section 3.1. According to the BIC, the short-run dynamics of the realized variance depends on five lags of the differenced rv_t series and three lags of the differenced vix_t^2 series, and the short-run dynamics of the implied variance depends on two lags of the differenced rv_t series, three lags of the differenced vix_t^2 series, and one lag of the differenced r_t series. Meanwhile, the return equation does not require the inclusion of any short-run lags. We refer to this particular specification as the CFVAR$_{0.4}(5)$ model below. The corresponding standard errors for the parameter estimates reported in the right column of the table are based on three thousand replications of a moving block bootstrap. Specifically, using the semiparametric estimates for d we first fractionally filter the two variance series. We then jointly resample blocks of the trivariate $I(0)$ vector $(\Delta^d rv_t,\Delta^d vix_t^2,r_t)'$, with the length of the blocks set to 3.5 months. We then apply the inverse fractional filter to the resampled variance series and finally reestimate the CFVAR model using the same SUR approach discussed in the text. With the exception of the

[11] Formal regularity conditions for the ϵ_t white noise process are spelled out in Johansen (2008a,b).

[12] The application of the fractional difference operator Δ^d to the returns might seemingly result in over-differencing. However, as shown in Appendix A.1 for the specific representation of the CFVAR model adopted here, the resulting return series is still $I(0)$.

[13] Note, the operator $(1-\Delta^d)$ involves only lags L^j, $j\geq 1$, with no L^0 term.

[14] We again truncate the fractional filter at the beginning of the sample, discarding the first weeks worth of observation, so that the estimation is based on the identical September 30, 2003 to December 30, 2011 sample period underlying the previously reported results in Tables 2 and 3.

Table 4
Co-fractional vector autoregression (CFVAR) model estimates.
The table reports seemingly unrelated regression (SUR) estimates of the $\text{CFVAR}_{0.40}(5)$ model,

$$\Delta^d z_t = \gamma(\delta'(1-\Delta^d)z_t+\rho)+\sum_{i=1}^{5}\Gamma_i\Delta^d(1-\Delta^d)^i z_t+\epsilon_t,$$

based on five-minute observations from September 30, 2003 to December 30, 2011. The fractional difference parameter is fixed at $d=0.40$. The reported standard errors (SE) for the parameter estimates are calculated from the bootstrap procedure discussed in the text.

Parameters	Estimates	SE
ρ'	$(-0.000 \quad -0.0000)$	$(0.111 \quad 0.0001)$
γ'	$\begin{pmatrix} 0.000473 & -0.00212 & 0.00542 \\ -0.00201 & -0.0325 & -1.089 \end{pmatrix}$	$\begin{pmatrix} 0.000078 & 0.00030 & 0.00124 \\ 0.00131 & 0.0065 & 0.010 \end{pmatrix}$
δ'	$\begin{pmatrix} -1.070 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$	$\begin{pmatrix} 0.202 & - & - \\ - & - & - \end{pmatrix}$
Γ_1	$\begin{pmatrix} 1.639 & -0.048 & 0 \\ 0.071 & 1.445 & 0.0076 \\ 0 & 0 & 0 \end{pmatrix}$	$\begin{pmatrix} 0.066 & 0.034 & - \\ 0.030 & 0.022 & 0.0027 \\ - & - & - \end{pmatrix}$
Γ_2	$\begin{pmatrix} -0.542 & 0.090 & 0 \\ -0.101 & -0.233 & 0 \\ 0 & 0 & 0 \end{pmatrix}$	$\begin{pmatrix} 0.112 & 0.073 & - \\ 0.033 & 0.043 & - \\ - & - & - \end{pmatrix}$
Γ_3	$\begin{pmatrix} -0.062 & -0.053 & 0 \\ 0 & -0.207 & 0 \\ 0 & 0 & 0 \end{pmatrix}$	$\begin{pmatrix} 0.049 & 0.040 & - \\ - & 0.030 & - \\ - & - & - \end{pmatrix}$
Γ_4	$\begin{pmatrix} 0.140 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$	$\begin{pmatrix} 0.104 & - & - \\ - & - & - \\ - & - & - \end{pmatrix}$
Γ_5	$\begin{pmatrix} -0.159 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$	$\begin{pmatrix} 0.046 & - & - \\ - & - & - \\ - & - & - \end{pmatrix}$

semiparametric first-step estimation of d, this bootstrap approach closely mimics that of Davidson (2002).

To gauge the model fit, Fig. 5 compares the spectra of the estimated $\text{CFVAR}_{0.4}(5)$ model with the sample periodograms for rv_t, vix_t^2, and r_t, as well as the variance difference $vd_t = vix_t^2 - rv_t$. The match between the model-implied and empirical spectra are exemplary. Particularly noteworthy, the CFVAR model correctly matches the slopes of the spectra near the origin that define the long-run behavior of the two variance series and the less persistent variance difference.

The cointegrating vector associated with the realized and risk-neutral variances is estimated to $(-1.070,1,0)$. This fully parametric CFVAR-based estimate for $\tilde{\delta}$ is even closer to unity (numerically) than the FDLS estimate discussed in Section 3.2. A simple t-test does not reject the hypothesis that $\tilde{\delta}=1$ either, while the absence of fractional cointegration, or $\tilde{\delta}=0$, is strongly rejected by the data. As such, our results are fully supportive of the notion of long-run unbiasedness in variances.[15]

To more clearly illuminate the dynamic dependencies implied by the $\text{CFVAR}_{0.4}(5)$ model, it is instructive to consider Impulse-Response Functions (IRF) associated with the shocks to the "permanent" and two transitory components defined within the model. By definition, the effect of the permanent shock decays hyperbolically and, therefore, persists over relatively long time periods. The two transitory shocks both decay at a fast exponential rate. The identification of these shocks does not depend a priori on any known relations between the volatility and return series.[16] However, given the estimates of γ and δ in Table 4, the three shocks have clear meanings. The first transitory shock, in particular, drives the wedge between vix_t^2 and rv_t, i.e., the variance difference, and contains the changes to the variance risk premium. The second transitory shock affects only the returns. The permanent shock is effectively a shock to rv_t that is unrelated to changes in the variance difference. This shock naturally also affects vix_t^2 through expectations of future variances.

Panel A in Fig. 6 shows the effect of the permanent shock on vix_t^2 and rv_t. The initial effect is largest for the vix_t^2, but this reverses after a couple of days when the impact on rv_t is the largest. Some discrepancy remains in the effect of the shock up until the 3.5-month horizon, but the overlapping nature of the rv_t series complicates the interpretation of this difference. After 3.5 months, the IRFs seem to merge and decay at a common hyperbolical rate. Even at the one-year horizon, however, the permanent shock still exercises a non-negligible effect on the rv_t and vix_t^2 series.

Panel B in Fig. 6 shows the IRFs for the first orthogonalized transitory shock associated with the variance difference. This shock fully dissipates after approximately 3.5 months. It essentially reflects a shock to the variance

[15] Comparable results based on coarser sampled monthly realized and options implied variances and semiparametric estimates for $\tilde{\delta}\approx 1$ have previously been reported by Bandi and Perron (2006), among others.

[16] Following Gonzalo and Granger (1995), the permanent and transitory shocks could be constructed mechanically from the ϵ_t's by multiplication with the matrix $G=[\gamma_\perp \delta]'$. The shocks are further orthogonalized using the method in Gonzalo and Ng (2001).

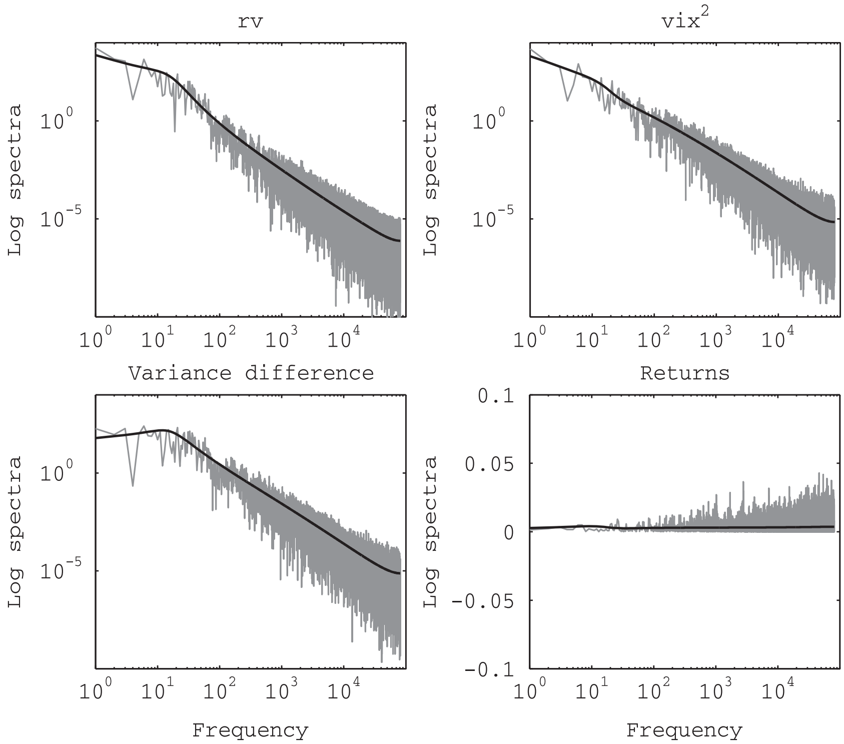

Fig. 5. Co-fractional vector autoregression (CFVAR) model implied spectra. The figure plots the CFVAR model implied spectra (solid lines) for the realized variance rv_t, the risk-neutral variance vix_t^2, the variance difference vd_t, and the returns r_t, along with their corresponding sample periodograms (grey lines). All of the estimates are based on five-minute observations from September 30, 2003 to December 30, 2011.

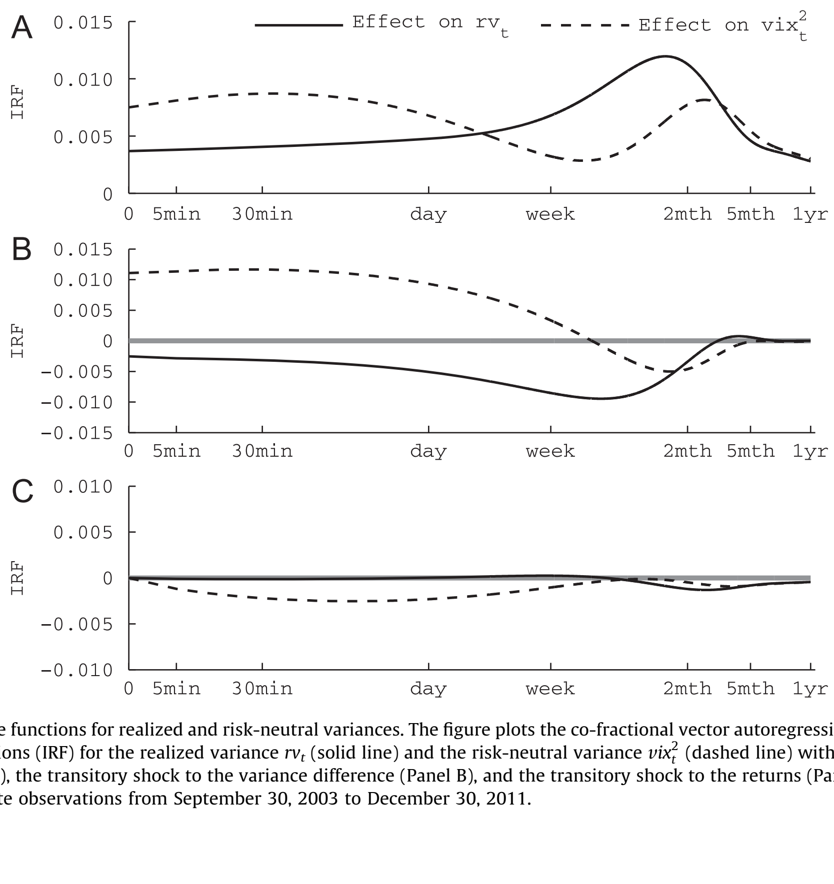

Fig. 6. Impulse response functions for realized and risk-neutral variances. The figure plots the co-fractional vector autoregression (CFVAR) model implied impulse response functions (IRF) for the realized variance rv_t (solid line) and the risk-neutral variance vix_t^2 (dashed line) with respect to the permanent variance shock (Panel A), the transitory shock to the variance difference (Panel B), and the transitory shock to the returns (Panel C). All of the estimates are based on five-minute observations from September 30, 2003 to December 30, 2011.

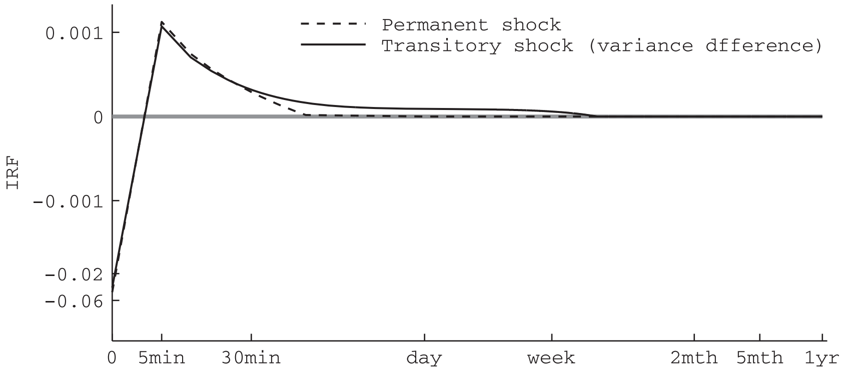

Fig. 7. Impulse response functions for returns. The figure plots the co-fractional vector autoregression (CFVAR) model implied impulse response functions (IRF) for the return r_t with respect to the permanent variance shock (dashed line) and the transitory shock to the variance difference (solid line). All of the estimates are based on five-minute observations from September 30, 2003 to December 30, 2011.

risk premium and a shock to the realized variance forecast error, or vrp_t and ξ_t, respectively, in the notation of Eq. (15). The first effect naturally increases vix_t^2, and the second effect decreases rv_t. The first effect is by far the largest over within-day horizons.

Panel C in the figure shows the effects of the second transitory return shock, from which the simultaneous correlations with the variance and variance difference shocks have been removed. This shock has no initial effect on either variance series. However, there is a negative impact on the vix^2 lasting for up to a week akin to a multi-period dynamic leverage effect, as previously discussed by Bollerslev, Litvinova, and Tauchen (2006).

In Fig. 7, we show the IRFs for the returns and the two variance shocks. The initial effect both shocks is to decrease the returns. This is entirely consistent with the widely documented leverage effect and negative contemporaneous correlations between returns and volatilities. Both of the shocks, in turn, result in an increase in future returns, as would be expected by a volatility feedback-type effect.

Taken as a whole, the IRFs suggest the potential for nontrivial return predictability through the joint CFVAR modeling of the returns and the two variance measures.

5.2. *Return predictability*

Until now we have relied on the forward monthly rv_t, or the realized variance over the subsequent month relative to time t. This ensures that rv_t is properly aligned with the future expected variance underlying the definition of vix_t^2. This also ties in more directly with the theoretical model in Bollerslev, Sizova, and Tauchen (2012) and the forward (expected) variance risk premium defined therein.

However, the "structural" CFVAR$_{0.4}$(5) model for $z_t = (rv_t, vix_t^2, r_t)'$ estimated above does not easily lend itself to out-of-sample forecasting, as rv_t is not known at time t. To circumvent this problem, we replace rv_t with $\tilde{rv}_t \equiv rv_{t-78\times 22}$ in the estimation of a CFVAR model for the $\tilde{z}_t \equiv (\tilde{rv}_t, vix_t^2, r_t)'$ vector process. The CFVAR$_{0.4}$(5) model for z_t formally implies a CFVAR$_{0.4}(p)$ model for $\tilde{z}_t$. We, therefore, fix $d=0.4$ and rely on the same BIC used above for determining the number of lags in each of the three equations in this new CFVAR$_{0.4}$(8) model for $\tilde{z}_t$.[17]

To convert the dynamic relations implied by the CFVAR model for $\tilde{z}_t$ into forecasts for the returns and corresponding predictive R^2's, it is convenient to represent the model in moving average form. In particular, let $e3' \equiv (0,0,1)$ so that $r_t = e3'\tilde{z}_t$. The implied infinite moving average representation for the returns can then be expressed as

$$r_t = e3' \sum_{j=0}^{\infty} \Phi_j \epsilon_{t-j}, \tag{23}$$

where the impulse responses matrices Φ_j follow the recursion

$$\Phi_0 = I,$$

$$\Phi_j = \sum_{i=0}^{j-1} \Xi_{j-i} \Phi_i, \tag{24}$$

for

$$\Xi_i = -I\theta_i^{(0)}(d) - (\gamma\delta')\theta_i^{(0)}(d) + \sum_{j=1}^{i} (-1)^j \Gamma_j \theta_i^{(j)}(d), \tag{25}$$

and the previously defined parameters of the fractional filter are given by $\theta_i^{(0)}(d) = (-1)^i \binom{d}{i}$ and $\theta_i^{(j)}(d) = \sum_{l=j-1}^{i-1} \theta_{i-l}^{(0)}(d)\theta_l^{(j-1)}(d)$, $j \geq 1$, respectively. These expressions readily allow for the calculation of h-step ahead return forecasts by simply equating all of the values of ϵ_{t+h-j} for $h-j > 0$ to zero in the corresponding expression for r_{t+h} in Eq. (23).

In practice, we are typically interested in the cumulative forecasts of the high-frequency returns over longer time intervals, as opposed to the multi-period forecasts of the high-frequency returns themselves. For illustration,

[17] The BIC implies that the short-run dynamics of the backward-looking realized variance depends on eight lags of the differenced $\tilde{rv}_t$ series and one lag of the differenced vix_t^2 series, while the short-run dynamics of the implied variance depends on three lags of the differenced vix_t^2 series and one lag of the differenced returns. The short-run dynamics of the returns depends on two lags of the differenced rv_t series. Additional details concerning these estimation results are available upon request.

consider the case of one-day returns. With 78 intraday returns per trading day, the continuously compounded daily return can be written as

$$r_t^{(\mathrm{day})} = \sum_{j=0}^{78-1} r_{t+j} = e3' \sum_{j=0}^{78-1} \sum_{i=0}^{\infty} \Phi_i \epsilon_{t+j-i}. \tag{26}$$

Going one step further, this expression for the daily return is naturally decomposed into an expected and an unexpected part,

$$r_t^{(\mathrm{day})} = e3' \sum_{j=0}^{78-1} \sum_{i=j+1}^{\infty} \Phi_i \epsilon_{t+j-i} + e3' \sum_{j=0}^{78-1} \sum_{i=0}^{j} \Phi_i \epsilon_{t+j-i}, \tag{27}$$

with the first term on the right-hand side corresponding to the former and the second term the latter. Consequently, the R^2 for the daily return implied by the CFVAR model can be conveniently expressed as

$$R^2_{\mathrm{day}} = \frac{\sum_{k=1}^{\infty} e3' (\sum_{j=0}^{78-1} \Phi_{k+j}) \Omega (\sum_{j=0}^{78-1} \Phi_{k+j})' e3}{\sum_{k=-(78-1)}^{\infty} e3' (\sum_{j=\max(0,-k)}^{78-1} \Phi_{k+j}) \Omega (\sum_{j=\max(0,-k)}^{78-1} \Phi_{k+j})' e3}. \tag{28}$$

Similar expressions for the R^2's associated with forecasting five-minute, hourly, weekly, and monthly returns are readily available by replacing 78 in the equation above with 1, 12, 390, and 1,716. For returns that are initially standardized by the diurnal pattern weights $\{1/w_1, \ldots, 1/w_{78}\}$ for $\{r_t, \ldots, r_{t+78-1}\}$ respectively, the impulse response functions Φ_{k+j} in Eq. (28) are replaced by $w_{j+1}\Phi_{k+j}$. The unconditional R^2_{day} is calculated by averaging over possible combinations of weights $\{w_1, \ldots, w_{78}\}$ for different times of the day.

The results obtained from evaluating the comparable expression in Eq. (28) for the five-minute, hourly, weekly, and monthly forecast horizons at the $\mathrm{CFVAR}_{0.4}(8)$ model estimates for $\tilde{z}_t$ are reported in the first row in Table 5.[18] As seen from the table, the CFVAR model for $\tilde{z}_t$ implies substantial R^2's of 1.683% and 6.354% at the weekly and monthly horizons, respectively. As such, these results further corroborate the empirical evidence pertaining to return predictability and a significant risk-return trade-off relation discussed above.

To help gauge where this predictability is coming from, we calculate the implied R^2's for three restricted versions of the $\mathrm{CFVAR}_{0.4}(8)$ model. In the first and second models, we restrict the returns to depend on the lags of $\tilde{rv}_t$ and vix_t^2, respectively, leaving the other dynamic dependencies in the $\mathrm{CFVAR}_{0.4}(8)$ model intact. The third restricted CFVAR model rules out any volatility feedback effects, so that the returns simply follow an autoregressive model. In each case, we reestimate the CFVAR model with the relevant restrictions imposed on the return equation and compute the corresponding VAR representation and implied R^2's. Further details concerning the calculation of the R^2's for the restricted CFVAR models are given in Appendix A.2.

The resulting R^2's are reported in rows 2 through 4 in Table 5. The models that include only the lagged $\tilde{rv}_t$ and vix_t^2 in the return equation result in almost no predictability at the daily horizon, but the R^2's increase to 1.390% and 4.789%, respectively, at the monthly horizon. By contrast, a simple autoregressive model for the returns, which does not include any of the variances in the return equation, performs comparatively well in predicting five-minute and hourly returns, but it has less forecasting power over longer horizons. By accounting for the joint dynamic dependencies in the returns and the two separate variance measures, the general $\mathrm{CFVAR}_{0.4}(8)$ model effectively combines the forecasting performance of the autoregressive model at the highest frequencies with that of the risk-based models at lower frequencies and, in turn, the predictability inherent in the variance risk premium.

To further underscore the gains afforded by jointly modeling the three series, the last three rows in Table 5 report the results from a set of simple univariate balanced predictive regressions, in which we regress the five-minute, hourly, daily, weekly, and monthly future returns on $\Delta^d \tilde{rv}_t$, $\Delta^d vix_t^2$, and $vp_t = vix_t^2 - \tilde{rv}_t$. These regressions directly mirror the non-predictive regressions of the contemporaneous returns on the same three predictor variables reported in Table 2. In parallel to the results for CFVAR models, all of the regressions are estimated with the standardized returns, while the reported R^2's are based on the raw returns. As expected, all of the R^2's are noticeably lower than those for the general $\mathrm{CFVAR}_{0.4}(8)$ model reported in the first row of the table.

6. Conclusion

We provide a detailed characterization of the dynamic dependencies and interrelatedness in aggregate stock market returns and volatilities using newly available high-frequency intraday data on both. The time series of actual realized volatilities and the market's risk-neutralized expectation thereof are both well described by long-memory fractional integrated processes. At the same time, the two volatility processes appear to be fractionally cointegrated and move in a one-to-one relation with one another in the long run. Using frequency domain inference procedures that allow us to focus on specific components of the spectra, we also uncover strong evidence for an otherwise elusive positive risk-return trade-off relation in the high-frequency data. Instead of a trade-off between returns and variances, however, the data clearly point to a trade-off between returns and the cointegrating relation between the two variance series. Moreover, we show that the strength of this relation varies importantly across frequencies and, as a result, simple risk-return regressions, as estimated in much of the existing literature, can easily give rise to misleading conclusions. Combining these results, we formulate and estimate a fractionally cointegrated VAR model for the high-frequency returns and two variance series that is able to accommodate all of these dependencies within a coherent joint modeling framework. Going one step further, we show how this high-frequency-based multivariate model implies nontrivial return predictability over longer monthly horizons.

[18] In carrying out the numerical calculations, we truncated the infinite sum in Eq. (28) at 100,000. Additional robustness checks using longer truncation lags and longer return horizons reveal the same basic conclusions.

Table 5
Multi-period return predictions.
The first row reports the predictive R^2's for five-minute, hourly, daily, weekly and monthly returns implied by the predictive $\text{CFVAR}_{0.40}(7)$ model discussed in the text. Rows 2 through 4 give the predictive R^2's for the restrictive co-fractional vector autoregression (CFVAR) models in which the return equation includes only lagged values of $\tilde{rv}_t$, vix_t^2, and the returns, respectively. The final three rows give the R^2's from simple univariate predictive return regressions. All of the regressions are estimated with the standardized returns, as discussed in the text, while all of the reported R^2's are calculated for the raw unstandardized returns. The five-minute observations underlying the estimates span the period from September 30, 2003 to December 30, 2011.

Model	$r_t^{(5\text{ min})}$	$r_t^{(\text{hour})}$	$r_t^{(\text{day})}$	$r_t^{(\text{week})}$	$r_t^{(\text{month})}$
$\text{CFVAR}_{0.40}(7)$	0.160%	0.150%	0.375%	1.638%	6.354%
CFVAR with $\tilde{rv}_t$ only	0.008%	0.050%	0.030%	0.222%	1.390%
CFVAR with vix_t^2 only	0.003%	0.073%	0.184%	1.135%	4.789%
CFVAR with no volatility feedback	0.137%	0.083%	0.013%	0.012%	0.856%
Regression with $\Delta^d \tilde{rv}_t$	0.000%	0.003%	0.039%	0.316%	0.870%
Regression with $\Delta^d vix_t^2$	0.014%	0.012%	0.000%	0.000%	0.069%
Regression with $vp_t = vix_t^2 - \tilde{rv}_t$	0.003%	0.026%	0.153%	0.592%	0.604%

The cointegrating relation between the two variance series uncovered in the paper is intimately related to the variance risk premium. The new class of stochastic volatility models proposed by Barndorff-Nielsen and Veraart (2013) provides a direct link between the variance risk premium and time-varying volatility-of-volatility. Our quantitative findings are also consistent with the qualitative implications from the stylized equilibrium model in Bollerslev, Sizova, and Tauchen (2012), in which the variance risk premium is linked to notions of aggregate economic uncertainty and time-varying equity risk premia. It would be interesting to further expand on these models to allow for a more structural explanation of the intricate cross-frequency empirical relations and dynamic dependencies in returns and risk-neutral and realized variances documented here.

Appendix A

A.1. CFVAR model solution for returns

Let $e1' \equiv (1,0,0)$, $e2' \equiv (0,1,0)$, and $e3' \equiv (0,0,1)$. The preliminary univariate estimates for the two variance series in Section 3 suggest that the first and the second equations of the $\text{CFVAR}_d(p)$, $e1'\Delta^d z_t \equiv \Delta^d rv_t$ and $e2'\Delta^d z_t \equiv \Delta^d vix_t^2$, respectively, are both $I(0)$. But the fractional filter Δ^d is also applied to the third equation $e3'\Delta^d z_t$ and the returns. This seemingly could result in over-differencing. However, following Theorem 8 of Johansen (2008b), if the conditions for inversion of the $\text{CFVAR}_d(p)$ are satisfied and $d < \frac{1}{2}$, then Eq. (20) has the solution

$$z_t = D\sum_{i=0}^{\infty}\theta_i^{(0)}(-d)L^i\epsilon_t + Y_t + \mu_t, \tag{29}$$

where $\theta_i^{(0)}(-d) = (-1)^i\binom{-d}{i}$ are the coefficients of the inverse fractional filter, μ_t is a function of the restricted constant ρ, Y_t is a stationary $I(0)$ series, and the parameter matrix D is defined by

$$D = \delta_\perp\left(\gamma_\perp'\left[I - \sum_{i=1}^{p}\Gamma_i\right]\delta_\perp\right)^{-1}\gamma_\perp', \tag{30}$$

where $\gamma_\perp$ and $\delta_\perp$ are 3×1 vectors such that $\gamma'\gamma_\perp = 0$ and $\delta'\delta_\perp = 0$, respectively. With γ and δ defined by Eqs. (22) and (21), the last row of D, therefore, has only zero elements. Consequently, the solution of the CFVAR model for the third equation and the return process reduces to

$$r_t = e3'Y_t + e3'\mu_t, \tag{31}$$

which is an $I(0)$ process, plus the initial contribution associated with ρ.

A.2. CFVAR restricted R^2's

For illustration, consider a daily forecast horizon. Let

$$z_t = \sum_{j=0}^{\infty}\tilde{\Phi}_j\tilde{\epsilon}_{t-j} \tag{32}$$

denote the moving average representation of the CFVAR model for z_t with the appropriate restrictions on the return equation imposed on the coefficients. Moving average coefficients are calculated using the recursions in Eq. (24). Define the polynomial

$$\Psi_1 + \Psi_2 L + \Psi_3 L^2 + \cdots = \left(\sum_{j=1}^{\infty}\sum_{i=0}^{78-1}\tilde{\Phi}_{j+i}L^j\right)\left(\sum_{j=0}^{\infty}\tilde{\Phi}_jL^j\right)^{-1}\left(\sum_{j=0}^{\infty}\Phi_jL^j\right). \tag{33}$$

The daily return forecast implied by the restricted CFVAR model is then given by $e3' \times \sum_{j=1}^{\infty}\Psi_j\epsilon_{t-j}$. The fraction of the cumulative daily returns that can be explained by the restricted model could, therefore, be computed as

$$\tilde{R}^2_{\text{day}} = \frac{\sum_{k=1}^{\infty} e3'\Psi_k\Omega(\sum_{j=0}^{78-1}\Phi_{k+j})'e3}{\sqrt{\sum_{k=1}^{\infty} e3'\Psi_k\Omega\Psi_k{}'e3 \times \sum_{k=-(78-1)}^{\infty} e3'\Sigma_{\Phi,k}\Omega\Sigma_{\Phi,k}{}'e3}}, \tag{34}$$

where $\Sigma_{\Phi,k} = \sum_{j=\max(0,-k)}^{78-1}\Phi_{k+j}$. Again, similar expressions for the five-minute, hourly, weekly, and monthly returns are readily available by replacing 78 in the formula above by the integer value corresponding to the relevant forecast horizon.

References

Andersen, T.G., Bollerslev, T., 1997a. Heterogeneous information arrivals and return volatility dynamics: uncovering the long-run in high frequency returns. Journal of Finance 52, 975–1005.

Andersen, T.G., Bollerslev, T., 1997b. Intraday periodicity and volatility persistence in financial markets. Journal of Empirical Finance 4, 115–158.

Andersen, T.G., Bollerslev, T., Diebold, F.X., Ebens, H., 2001. The distribution of realized stock return volatility. Journal of Financial Economics 61, 43–76.

Andersen, T.G., Bollerslev, T., Diebold, F.X., Labys, P., 2001. The distribution of realized exchange rate volatility. Journal of the American Statistical Association 96, 42–55.

Andersen, T.G., Bollerslev, T., Diebold, F.X., Labys, P., 2003. Modeling and forecasting realized volatility. Econometrica 71, 579–625.

Backus, D.K., Gregory, A.W., 1993. Theoretical relations between risk premiums and conditional variances. Journal of Business and Economic Statistics 11, 177–185.

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

Baillie, R.T., Bollerslev, T., 2000. The forward premium anomaly is not as bad as you think. Journal of International Money and Finance 19, 471–488.

Baillie, R.T., Bollerslev, T., Mikkelsen, H.O., 1996. Fractionally integrated generalized autoregressive conditional heteroskedasticity. Journal of Econometrics 74, 3–30.

Bakshi, G., Kapadia, N., 2003. Delta-hedged gains and the negative market volatility risk premium. Review of Financial Studies 16, 527–566.

Bali, T.G., Peng, L., 2006. Is there a risk-return trade-off? Evidence from high-frequency data. Journal of Applied Econometrics 21, 1169–1198.

Bandi, F.M., Perron, B., 2006. Long memory and the relation between implied and realized volatility. Journal of Financial Econometrics 4, 636–670.

Barndorff-Nielsen, O.E., Hansen, P.R., Lunde, A., Shephard, N., 2008. Designing realised kernels to measure the ex-post variation of equity prices in the presence of noise. Econometrica 76, 1481–1536.

Barndorff-Nielsen, O.E., Shephard, N., 2002. Econometric analysis of realised volatility and its use in estimating stochastic volatility models. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 64, 253–280.

Barndorff-Nielsen, O.E., Veraart, A.E.D., 2013. Stochastic volatility of volatility and variance risk premia. Journal of Financial Econometrics 11, 1–46.

Baxter, M., King, R.G., 1999. Measuring business cycles: approximate band-pass filters for economic time series. Review of Economics and Statistics 81, 575–593.

Black, F., 1976. Studies of stock price volatility changes. In: Proceedings of the 1976 Meetings of the American Statistical Association, Business and Economics Statistics Section, pp. 177–181.

Bollerslev, T., Gibson, M., Zhou, H., 2011. Dynamic estimation of volatility risk premia and investor risk aversion from option-implied and realized volatilities. Journal of Econometrics 160, 235–245.

Bollerslev, T., Litvinova, J., Tauchen, G., 2006. Leverage and volatility feedback effects in high-frequency data. Journal of Financial Econometrics 4, 353–384.

Bollerslev, T., Sizova, N., Tauchen, G., 2012. Volatility in equilibrium: asymmetries and dynamic dependencies. Review of Finance 16, 31–80.

Bollerslev, T., Tauchen, G., Zhou, H., 2009. Expected stock returns and variance risk premia. Review of Financial Studies 22, 4463–4492.

Campbell, J.Y., Hentschell, L., 1992. No news is good news: an asymmetric model of changing volatility in stock returns. Journal of Financial Economics 31, 281–318.

Carr, P., Wu, L., 2009. Variance risk premiums. Review of Financial Studies 22, 1311–1341.

Chernov, M., 2007. On the role of risk premia in volatility forecasting. Journal of Business and Economic Statistics 25, 411–426.

Christensen, B.J., Nielsen, M.Ø., 2006. Asymptotic normality of narrow-band least squares in the stationary fractional cointegration model and volatility forecasting. Journal of Econometrics 133, 343–371.

Christensen, B.J., Nielsen, M.Ø., 2007. The effect of long memory in volatility on stock market fluctuations. Review of Economics and Statistics 89, 684–700.

Comte, F., Renault, E., 1998. Long memory in continuous-time stochastic volatility models. Mathematical Finance 8, 291–323.

Corsi, F., 2009. A simple approximate long-memory model of realized volatility. Journal of Financial Econometrics 7, 174–196.

Davidson, J., 2002. A model of fractional cointegration, and tests for cointegration using the bootstrap. Journal of Econometrics 110, 187–212.

Ding, Z., Granger, C.W.J., Engle, R.F., 1993. A long memory property of stock market returns and a new model. Journal of Empirical Finance 1, 83–106.

Engle, R.F., Lilien, D.M., Robins, R.P., 1987. Estimating time varying risk premia in the term structure: the ARCH-M model. Econometrica 55, 391–407.

French, K.R., Schwert, G.W., Stambaugh, R.F., 1987. Expected stock returns and volatility. Journal of Financial Economics 19, 3–29.

Geweke, J., Porter-Hudak, S., 1983. The estimation and application of long memory time series models. Journal of Time Series Analysis 4, 221–237.

Glosten, L.R., Jagannathan, R., Runkle, D.E., 1993. On the relation between the expected value and the volatility of the nominal excess return on stocks. Journal of Finance 48, 1779–1801.

Gonzalo, J., Granger, C.W.J., 1995. Estimation of common long-memory components in cointegrated systems. Journal of Business and Economic Statistics 13, 27–35.

Gonzalo, J., Ng, S., 2001. A systematic framework for analyzing the dynamic effects of permanent and transitory shocks. Journal of Economic Dynamics and Control 25, 1527–1546.

Granger, C.W.J., 1986. Developments in the study of cointegrated economic variables. Oxford Bulletin of Economics and Statistics 48, 213–228.

Granger, C.W.J., Hatanaka, M., 1964. Spectral Analysis of Economic Time Series. Princeton University Press, Princeton, NJ, USA.

Harris, L., 1986. A transaction data study of weekly and intradaily patterns in stock returns. Journal of Financial Economics 16, 99–117.

Hassler, J., Lundvik, P., Persson, T., Söderlind, P., 1994. The Swedish business cycle: stylized facts over 130 years. In: Bergström, V., Vredin, A. (Eds.), Measuring and Interpreting Business Cycles, Clarendon Press, Oxford, UK, pp. 11–123.

Jacod, J., Li, Y., Mykland, P.A., Podolskij, M., Vetter, M., 2009. Microstructure noise in the continuous time case: the pre-averaging approach. Stochastic Processes and Their Applications 119, 2249–2276.

Johansen, S., 2008a. A representation theory for a class of vector autoregressive models for fractional processes. Econometric Theory 24, 651–676.

Johansen, S., 2008b. Representation of cointegrated autoregressive processes with application to fractional processes. Econometric Reviews 28, 121–145.

Künsch, H., 1987. Statistical aspects of self-similar processes. In: Prohorov, Y.V., Sazonov, V.V. (Eds.), Proceedings of the First World Congress of the Bernoulli Society. VNU Science Press, Utrecht, The Netherlands, pp. 67–74.

Maynard, A., Smallwood, A., Wohar, M.E., 2013. Long memory regressors and predictive regressions: a two-stage rebalancing approach. Econometric Reviews 32, 318–360.

Merton, R.C., 1980. On estimating the expected return on the market: an exploratory investigation. Journal of Financial Economics 8, 323–361.

Nelson, D.B., 1991. Conditional heteroskedasticity in asset returns: a new approach. Econometrica 59, 347–370.

Ohanissian, A., Russell, J.R., Tsay, R.S., 2009. True or spurious long memory? A new test. Journal of Business and Economic Statistics 26, 161–175.

Robinson, P.M., 1994. Semiparametric analysis of long-memory time series. Annals of Statistics 22, 515–539.

Robinson, P.M., 1995. Log-periodogram regression of time series with long range dependence. Annals of Statistics 23, 1048–1072.

Robinson, P.M., Marinucci, D., 2003. Semiparametric frequency domain analysis of fractional cointegration. In: Robinson, P.M. (Ed.), Time Series with Long Memory, Oxford University Press, New York, NY, pp. 335–373.

Rossi, A., Timmermann, A., 2012. What is the shape of the risk-return relation? Unpublished working paper. University of California, Rady School of Management, San Diego, CA.

Shimotsu, K., Phillips, P.C.B., 2006. Local Whittle estimation of fractional integration and some of its variants. Journal of Econometrics 130, 209–233.

Sowell, F.B., 1992. Modeling long run behavior with the fractional ARIMA model. Journal of Monetary Economics 29, 277–302.

Zhang, L., Mykland, P.A., Aït-Sahalia, Y., 2005. A tale of two time scales: determining integrated volatility with noisy high-frequency data. Journal of the American Statistical Association 100, 1394–1411.