

On the macroeconomic determinants of long-term volatilities and correlations in U.S. stock and crude oil markets*

Christian Conrad[†], Karin Loch[‡], and Daniel Rittler[§]

Heidelberg University, Germany

First version: March 13, 2012

This version: August 1, 2013

Abstract

Using a modified DCC-MIDAS specification, we endogenize the long-term correlation between crude oil and stock price returns with respect to the stance of the U.S. macroeconomy. We find that variables which contain information on current and future economic activity are helpful predictors for changes in the oil-stock correlation. For the period 1993-2011 there is strong evidence for a counter cyclical behavior of the long-term correlation. For prolonged periods with strong growth above trend our model predicts a negative long-term correlation, while before and during recessions the sign changes and remains positive throughout the economic recovery.

Keywords: Oil-stock relationship, long-term volatility, long-term correlation, GARCH-MIDAS, DCC-MIDAS

JEL Classification: C32, C58, Q43

*We would like to thank Thomas Eife, Lutz Kilian, Sandra Schmidt, and Timo Teräsvirta for valuable comments and suggestions. Christian Conrad gratefully acknowledges financial support from the “Juniorprofessorenprogramm” of the state of Baden-Württemberg.

[†]Christian Conrad, Department of Economics, Heidelberg University, Bergheimer Strasse 58, 69115 Heidelberg, Germany, E-mail: christian.conrad@awi.uni-heidelberg.de; Phone: +49/6221/54/3173.

[‡]Karin Loch, Department of Economics, Heidelberg University, Bergheimer Strasse 58, 69115 Heidelberg, Germany, E-mail: karin.loch@awi.uni-heidelberg.de; Phone: +49/6221/54/2908.

[§]Daniel Rittler, E-mail: daniel.rittler@web.de

1 Introduction

In this article, we revisit the oil-stock market relationship by analyzing the macroeconomic determinants of the long-term correlation between daily U.S. stock market and crude oil price returns. Recently, Kilian and Park (2009) have shown that on average 22% of the variation in U.S. stock returns in the period 1975–2006 can be explained by oil price shocks. However, whether an oil price shock drives oil and stock prices in the same or in opposite directions crucially depends on the type of the underlying shock. While oil price increases due to precautionary demand have a negative effect on stock prices, demand driven oil price shocks lead to increasing stock prices. Based on these insights, Kilian and Park (2009) argue that the time-varying sign in rolling oil-stock correlations reflects changes in the relative importance of different demand and supply shocks in the oil market.

While Kilian and Park (2009) investigate the oil-stock relationship using monthly data, our purpose is to analyze the correlation between oil and stock returns at a daily frequency. More specifically, we use a novel MIXed Data Sampling (MIDAS) approach to link the smooth component of daily return correlations to changes in monthly U.S. macroeconomic variables. While there is a growing literature on the endogeneity of monthly or quarterly oil prices with respect to U.S. and global macroeconomic conditions (Barsky and Kilian, 2004; Kilian, 2008, 2009), our contribution is to provide first evidence on the link between U.S. economic activity and the daily oil-stock correlation.¹

Our econometric specification is based on the Dynamic Conditional Correlation - MIDAS (DCC-MIDAS) model proposed in Colacito et al. (2011). The DCC-MIDAS combines the Engle (2002) DCC specification with the GARCH-MIDAS framework of Engle et al. (2012). The latter framework extends the simple GARCH specification by modeling volatility as consisting of a short-term and a long-term component. Most importantly, the long-term component is specified as a function of the macroeconomic environment. In the original DCC specification with correlation targeting each quasi-correlation follows a ‘GARCH type’ process, which is mean-reverting to the unconditional correlation of the volatility-adjusted residuals. The basic idea of Colacito et al. (2011) is to replace this unconditional correlation with a slowly time-varying long-term component. The quasi-correlation then fluctuates around this long-run trend. Hence, the new specification can be

¹In the following, we refer to the correlation between oil and stock returns simply as the oil-stock correlation.

considered as a two-component model for the dynamic correlations. Colacito et al. (2011) model the long-term component as a weighted sum of the lagged monthly realized correlations between the volatility-adjusted residuals.

Using the GARCH-MIDAS framework, we first analyze whether the long-term oil market volatility is related to the U.S. macroeconomy and whether oil and stock volatility respond to the same macroeconomic information. We then extend the DCC-MIDAS model by directly incorporating information on the macroeconomic development in the long-term correlation component, i.e. we replace the realized correlations by monthly macroeconomic variables. Since the macroeconomic variables – unlike the realized correlations – are not restricted to the minus one to plus one interval, we suggest a new specification for the long-term component. Similar to Christodouklakis and Satchell (2002), we assume that the Fisher- z transformation of the long-term component can be written as a linear function of the weighted lagged macroeconomic variables. The weights are again determined using the MIDAS approach. We refer to this new specification, which includes a macroeconomic explanatory variable, as the DCC-MIDAS-X model.

Our results can be summarized as follows. First, we find that the movements in long-term oil market volatility can be well predicted by various measures of U.S. macroeconomic activity. Our empirical results provide convincing evidence for a counter cyclical relationship between oil market volatility and variables which either describe the current stance of the economy, e.g. industrial production, or provide forward looking information about the future state of the economy, e.g. the leading index for the U.S. Current and expected increases (decreases) in economic activity clearly anticipate downswings (upswings) in long-term oil volatility. While the notion that there is reverse causality from macroeconomic variables to the level of the oil price (see, e.g., Barsky and Kilian, 2004; Kilian, 2008, 2009) is now widely accepted, our result adds a new dimension by establishing a link between U.S. macroeconomic variables and the volatility of oil price returns. Interestingly, we also find that long-term oil and stock market volatility respond to the same macroeconomic information.

Second, our empirical results show that changes in the long-term oil-stock correlation can be anticipated by the same macroeconomic factors that affect the long-term volatilities. We provide strong evidence for a counter cyclical behavior of the long-term oil-stock correlation. Phases with positive long-term oil-stock correlations correspond to values of the macroeconomic factors which either indicate recessions or the beginning of expansions with growth still below or at trend. On the other hand, a negative long-run correlation

emerges when the macroeconomic variables signal strong growth above trend. Clearly, the positive correlation during recessions is driven by the simultaneous drop in oil and stock prices. The economic recovery during the early phase of an expansion then leads to increasing oil prices due to higher demand as well as to rising stock prices because of the improved outlook for corporate cash flows. The combination of these two effects causes the long-run oil-stock correlation to remain positive. This interpretation squares with the findings in Kilian and Park (2009) regarding the positive short-run effect on oil and stock prices of an unexpected increase in global demand. Finally, during boom phases with strong growth above trend both the further increasing oil prices as well as the expectation of rising interest rates have a depressing effect on the stock market. Hence, for these periods our model predicts a decreasing or negative long-term correlation.

Third, the long-term correlation component can be interpreted as the predicted or expected correlation given a certain state of the economy. Since the macroeconomic variables that drive the long-term component represent aggregate demand, the deviations of the short-term from the long-term component should be driven by other factors related to the stock and/or the oil market. Typical examples for the oil market would be either oil specific, i.e. precautionary, demand shocks or supply shocks. The behavior of the short-term component during the second Gulf War in 2003 (see Figure 3) is in line with this interpretation. However, the fact that various measures of macroeconomic activity lead to a convincing and coherent fit of the long-term correlation suggests that aggregate demand is the most important factor for the oil-stock relationship. This interpretation is very much in line with the view that – in contrast to the 1970s when supply shocks were likely to be predominant – oil prices have been mainly driven by high global aggregate demand since the mid-1990s (see Hamilton, 2008; Kilian, 2009; Kilian and Murphy, 2013).²

Fourth, the fact that the sign of the oil-stock correlation critically depends on the state of the economy reinforces Kilian and Park’s (2009) argument that simple regressions of stock returns on oil price changes can be very misleading. This point may well explain the conflicting empirical evidence on the oil-stock relationship in Jones and Kaul (1996), Wei (2003), Nandha and Faff (2008), Miller and Ratti (2009) and others.

Fifth, we show that the volatility and correlation predictions from the various DCC-MIDAS-X specifications significantly outperform the ones from the simple DCC model.

²Although we focus on economic activity measures for the U.S. only, while the oil price is driven by global demand, our approach may still be informative to the extent that changes in U.S. real activity are correlated with changes in global real activity.

Hence, the explicit modeling of the long-term correlation component may be very beneficial for portfolio choice, hedging decisions or risk management.

The remainder of the article is organized as follows. Section 2 reviews the related literature while Section 3 discusses the GARCH-MIDAS and DCC-MIDAS models. The data and empirical results are presented in Sections 4 and 5. In Section 6 we evaluate the forecasting performance of the different models and Section 7 concludes the article.

2 Related literature

Our analysis is based on two strands of literature. The first one is concerned with the modeling of long-term movements in volatilities and correlations, the second one with the relationship between oil and stock prices and macroeconomic conditions.

The idea of having short- and long-term component models of volatilities dates back to Ding and Granger (1996) and Engle and Lee (1999). In their specifications, both components follow ‘GARCH-type’ processes but with different degrees of persistence. Similarly, Davidson (2004) proposed the HYGARCH specification, which can be considered a two-component model with the short-term component being a GARCH process while the long-term component follows a FIGARCH process (see also Conrad, 2010). While these specifications allow one to separate the two volatility components, the unconditional variance is still assumed to be constant over time. Engle and Rangel (2008) and Engle et al. (2012) relax this assumption and propose specifications in which the long-term component can be considered a time-varying unconditional variance. While in the Engle and Rangel (2008) Spline-GARCH model both components fluctuate at the same frequency, in Engle et al. (2012) it is assumed that the long-term component evolves at a lower frequency than the short-term component. Using the MIDAS framework of Ghysels et al. (2005, 2007), they directly relate the long-term component to the evolution of macroeconomic time series such as industrial production or inflation. In line with the earlier findings in Schwert (1989), the GARCH-MIDAS model provides strong evidence for a counter cyclical behavior of financial volatility. Recently, Conrad and Loch (2012) extended the analysis of Engle et al. (2012) by using a broader set of macroeconomic variables including leading indicators and expectations data from the Survey of Professional Forecasters. The DCC-MIDAS model proposed in Colacito et al. (2011) simply extends the two-component idea from volatilities to correlations. However, instead of relating the long-term correlation directly to its potential macroeconomic sources, Colacito

et al. (2011) only consider lagged realized correlations as explanatory variables.

Since the seminal articles of Hamilton (1983, 1985, 2003) exogenous oil supply shocks were suspected to be causal for recessions and periods of low economic growth. Based on this presumption, several empirical studies have analyzed the relationship between oil prices and stock market returns. While Jones and Kaul (1996) or Nandha and Faff (2008) indeed find that oil price increases negatively affect stock prices, Huang et al. (1996) and Wei (2003) cannot establish a significant relationship. Recently, Miller and Ratti (2009) provide evidence for a time-varying relationship. For the period after 1999 they even report a positive connection. Hence, the empirical evidence is far from being uncontroversial. Kilian and Park (2009) provide two explanations for the conflicting results. First, there is convincing evidence for reverse causality from the U.S. economy to the oil price (see also Kilian, 2009, and Alquist et al., 2013). Thus, stock and oil price changes may be induced by the same macroeconomic factors and, hence, regressions of stock returns on oil price changes can be misleading due to endogeneity. Second, Kilian and Park (2009) argue that the sign of the effect of an oil price increase on the stock market depends on the type of the underlying shock and, hence, may change over time. While shocks due to an unanticipated economic expansion may have a positive impact, shocks related to precautionary demand, for example, are likely to have a negative impact. For several oil-exporting and oil-importing countries Filis et al. (2011) confirm that the oil-stock correlation is indeed time-varying. Although they informally relate phases of positive or negative correlations to demand and supply shocks, their simple DCC-GARCH model does not explicitly incorporate information on the state of the economy. In particular, their model does not allow one to forecast changes in correlations in response to changes in the macro environment.

3 The DCC-MIDAS model

In this section, we develop the econometric framework to analyze the impact of macroeconomic variables on long-term volatility and correlations. We consider the bivariate vector of asset returns $\mathbf{r}_t = (r_{1,t}, r_{2,t})'$, where $r_{1,t}$ refers to the stock and $r_{2,t}$ to the oil returns, and denote by $\mathcal{F}_{t-1} = \sigma(\mathbf{r}_{t-1}, \mathbf{r}_{t-2}, \dots)$ the σ -field generated by the information available through time $t - 1$. Let $\mathbf{E}[\mathbf{r}_t | \mathcal{F}_{t-1}] = \boldsymbol{\mu}_t = (\mu_{1,t}, \mu_{2,t})'$ and define the vector of residuals $\mathbf{r}_t - \boldsymbol{\mu}_t = \boldsymbol{\varepsilon}_t = (\varepsilon_{1,t}, \varepsilon_{2,t})'$. We assume that conditional on \mathcal{F}_{t-1} the residuals are normally distributed with $\mathbf{Var}[\boldsymbol{\varepsilon}_t | \mathcal{F}_{t-1}] = \mathbf{H}_t$, i.e. $\boldsymbol{\varepsilon}_t | \mathcal{F}_{t-1} \sim \mathcal{N}(\mathbf{0}, \mathbf{H}_t)$. Following Engle (2002), we

decompose the conditional covariance matrix into $\mathbf{H}_t = \mathbf{D}_t \mathbf{R}_t \mathbf{D}_t$ where

$$\mathbf{R}_t = \begin{pmatrix} 1 & \rho_{12,t} \\ \rho_{12,t} & 1 \end{pmatrix} \quad \text{and} \quad \mathbf{D}_t = \begin{pmatrix} h_{1,t}^{1/2} & 0 \\ 0 & h_{2,t}^{1/2} \end{pmatrix}. \quad (1)$$

Finally, we define the standardized residuals $\boldsymbol{\eta}_t = (\eta_{1,t}, \eta_{2,t})'$ as $\boldsymbol{\eta}_t = \mathbf{D}_t^{-1} \boldsymbol{\varepsilon}_t$. Note that $\mathbf{Var}[\boldsymbol{\eta}_t | \mathcal{F}_{t-1}] = \mathbf{R}_t$. The DCC framework allows us to separately model the conditional variances and the conditional correlations.

3.1 Conditional variances

To capture the impact of macroeconomic variables on return volatility, we adopt the GARCH-MIDAS framework of Engle et al. (2012). We assume a multiplicative component model for each conditional variance, i.e. we specify $h_{i,t} = g_{i,t} m_{i,\tau}$, where $g_{i,t}$ is the short-run and $m_{i,\tau}$ the long-run component. While the transitory volatility component changes at the daily frequency t , the long-run component changes at the monthly frequency τ only. We denote $N^{(\tau)}$ as the number of days within month τ . Specifically, we assume that the short-run volatility component follows a mean-reverting unit GARCH(1,1) process

$$g_{i,t} = (1 - \alpha_i - \beta_i) + \alpha_i \frac{(r_{i,t-1} - \mu_{i,t-1})^2}{m_{i,\tau}} + \beta_i g_{i,t-1}, \quad (2)$$

with $\alpha_i > 0$, $\beta_i \geq 0$, and $\alpha_i + \beta_i < 1$. The long-term component is modeled as a slowly varying function of exogenous variables X_τ using the MIDAS specification

$$\log(m_{i,\tau}) = m_i + \theta_i \sum_{k=1}^{K_v} \varphi_k(\omega_i) X_{\tau-k}, \quad (3)$$

where the log transformation guarantees the non-negativity of the conditional variances when the exogenous variables can take negative values. The X_τ will be monthly macroeconomic variables. For the weighting scheme, we follow Engle et al. (2012) and adopt a restricted beta weighting scheme where the weights are computed according to

$$\varphi_k(\omega_i) = \frac{(1 - k/K_v)^{\omega_i - 1}}{\sum_{l=1}^{K_v} (1 - l/K_v)^{\omega_i - 1}}, \quad k = 1, \dots, K_v. \quad (4)$$

For all $\omega_i > 1$, the weighting scheme guarantees a decaying pattern, where the rate of decay is determined by ω_i . Large (small) values of ω_i generate a rapidly (slowly) decaying pattern. The maximum lag length K_v is determined by an information criterion. By construction, the $\varphi_k(\omega_i)$ are nonnegative and sum to one.

In the following, we will refer to the component model with explanatory variables as GARCH-MIDAS-X. Finally, note that when $\theta_i = 0$ the long-run component is simply a constant and, hence, $h_{i,t}$ follows a stationary GARCH(1,1) process with constant unconditional variance.

3.2 Conditional correlations

The DCC-MIDAS specification proposed by Colacito et al. (2011) provides a natural extension of the GARCH-MIDAS model to dynamic correlations. We first decompose the conditional correlation matrix as $\mathbf{R}_t = \text{diag}\{\mathbf{Q}_t\}^{-1/2}\mathbf{Q}_t\text{diag}\{\mathbf{Q}_t\}^{-1/2}$, with $\mathbf{Q}_t = [q_{ij,t}]_{i,j=1,2}$, and specify the quasi-correlations as

$$\mathbf{Q}_t = (1 - a - b)\bar{\mathbf{R}}_t + a\boldsymbol{\eta}_{t-1}\boldsymbol{\eta}_{t-1}' + b\mathbf{Q}_{t-1}, \quad (5)$$

with $a > 0$, $b \geq 0$, and $a + b < 1$. In the Engle (2002) DCC model with correlation targeting the matrix $\bar{\mathbf{R}}_t$ does not depend on time and equals the empirical correlation matrix of $\boldsymbol{\eta}_t$, i.e. has ones on the main diagonal while the off-diagonal elements are given by $\bar{\rho}_{12} = T^{-1} \sum_{t=1}^T \eta_{1,t}\eta_{2,t}$. In contrast, in the DCC-MIDAS framework the off-diagonal elements are the long-term correlations $\bar{\rho}_{12,\tau}$. As in the GARCH-MIDAS equation the long-term correlation component does not vary at the daily frequency t but at the lower frequency τ . That is, the short-run quasi-correlations fluctuate around the time-varying long-run correlations:

$$q_{12,t} = \bar{\rho}_{12,\tau} + a(\eta_{1,t-1}\eta_{2,t-1} - \bar{\rho}_{12,\tau}) + b(q_{12,t-1} - \bar{\rho}_{12,\tau}). \quad (6)$$

Colacito et al. (2011) assume that $\bar{\rho}_{12,\tau}$ can be expressed as a weighted average of the K_c past realized correlations RC_τ :

$$\bar{\rho}_{12,\tau} = \sum_{k=1}^{K_c} \varphi_k(\omega_{12}) RC_{\tau-k}, \quad (7)$$

with

$$RC_\tau = \frac{\sum_{t=N_{\tau-1}+1}^{N_\tau} \eta_{1,t}\eta_{2,t}}{\sqrt{\sum_{t=N_{\tau-1}+1}^{N_\tau} \eta_{1,t}^2 \sum_{t=N_{\tau-1}+1}^{N_\tau} \eta_{2,t}^2}}, \quad (8)$$

where $N_\tau = \sum_{i=1}^\tau N^{(i)}$ and $N_0 = 0$. The weights are again given by equation (4) with ω_i and K_v replaced by ω_{12} and K_c , respectively. Since the weights $\varphi_k(\omega_{12})$ sum up to one and the RC_τ are correlations, the long-run correlation will itself lie within the $[-1, +1]$ interval.

We extend the DCC-MIDAS model by directly incorporating information on the macroeconomic development in the long-run component. Similarly as in the GARCH MIDAS setting – where the specification for $m_{i,\tau}$ has to ensure the non-negativity of the long-term volatility – our specification has to ensure that the long-run correlation lies within the $[-1, +1]$ interval although the explanatory variables do not. We follow Christodoulakis and Satchell (2002) and use the Fisher- z transformation of the correlation coefficient, i.e. we assume that

$$\bar{\rho}_{12,\tau} = \frac{\exp(2z_{12,\tau}) - 1}{\exp(2z_{12,\tau}) + 1}, \quad (9)$$

with

$$z_{12,\tau} = m_{12} + \theta_{12} \sum_{k=1}^{K_c} \varphi_k(\omega_{12}) X_{\tau-k}, \quad (10)$$

where X_τ denotes either a macroeconomic variable or a realized correlation. Note that in our non-linear specification, from θ we can only infer the sign but not directly the marginal effect of a macroeconomic variable on the long-term correlation.

Finally, in the DCC-MIDAS model - as in the standard DCC model - the short-run correlations are obtained by rescaling, i.e. $\rho_{12,t} = q_{12,t} / \sqrt{q_{11,t} q_{22,t}}$. In the subsequent analysis we refer to the specifications with either macroeconomic explanatory variables or the realized correlations as DCC-MIDAS-X or DCC-MIDAS-RC models, respectively.

3.3 Estimation

Following Engle (2002) and Colacito et al. (2011) the model parameters can be estimated using a two-step procedure. This is feasible because the log likelihood function to be maximized

$$\mathcal{L} = - \sum_{t=1}^T (2\log(2\pi) + 2\log(|\mathbf{D}_t|) + \boldsymbol{\varepsilon}_t' \mathbf{D}_t^{-2} \boldsymbol{\varepsilon}_t) - \sum_{t=1}^T (\log(|\mathbf{R}_t|) + \boldsymbol{\eta}_t' \mathbf{R}_t^{-1} \boldsymbol{\eta}_t - \boldsymbol{\eta}_t' \boldsymbol{\eta}_t) \quad (11)$$

can be separated into two parts. The first sum in equation (11) contains the data and the variance parameters while the second sum depends on the volatility-adjusted residuals and the correlation parameters. Hence, in the first step we estimate the GARCH-MIDAS parameters individually for each return series and use the estimated volatility-adjusted residuals in the second step to obtain the correlation parameters.

4 Data

We combine daily U.S. stock market and crude oil price data with monthly observations on the macroeconomic variables. While the stock series was obtained from the Kenneth R. French data library, the oil prices and the macroeconomic data are taken from the FRED database at the Federal Reserve Bank of St. Louis. Our data covers the period from January 1993 to November 2011.

4.1 Oil and stock market data

For the stock series, we employ the daily returns on the CRSP value-weighted portfolio, which is based on all NYSE, AMEX and NASDAQ stocks and can be considered the best available proxy for ‘the stock market’. As in Kilian and Vega (2011), the oil price returns are constructed from the West Texas Intermediate (WTI) crude oil spot price. Panel A of Table 1 provides summary statistics for the two return series.³ While the sample mean of the returns is positive for both markets, the table provides first evidence for stronger fluctuations in oil returns than in stock market returns. The annualized unconditional standard deviation of the oil price returns is 39.21% and, hence, considerably higher than the 19.53% of the CRSP returns. Finally, the unconditional correlation between oil and stock returns is 0.14.

Table 1 about here

4.2 Macroeconomic data

We divide the monthly macroeconomic data into two categories: those which measure the current stance of the economy and forward looking indicators.⁴ The first category contains the following variables: industrial production (IP), nonfarm payrolls (NFP), and the unemployment rate (UR). The forward looking indicators are the national activity

³As alternative measures for the stock market we also considered the S&P 500 as well as the DJIA. Similarly, we employed the Brent instead of the WTI crude oil price. All the subsequent results were robust to these changes in the variables.

⁴In a previous version of this paper, we also included CPI and PPI based inflation measures. However, for our sample we found neither a significant effect of inflation on long-term stock and oil market volatility nor on the long-term oil-stock correlation.

index (NAI)⁵ and the leading index (LI)⁶ for the U.S. They are supposed to reflect the role of market participants' expectations concerning the future economic development.

For the variables IP and NFP we compute month-to-month growth rates according to $100 \cdot [\ln(X_t) - \ln(X_{t-1})]$, while in case of UR we use month-to-month changes. The NAI and LI are included in levels. Panel B of Table 1 provides the summary statistics for the macroeconomic data and Figure 1 shows the dynamics of the macroeconomic variables. Note that by construction the GARCH- and DCC-MIDAS models need additional lags of the explanatory variables at the beginning of the sample. Since we shall include three MIDAS lag years in the filter, we report descriptive statistics and figures for the macroeconomic variables for the period from January 1990 to November 2011.

Figure 1 about here

5 Empirical results

We first present the estimation results for the GARCH-MIDAS models that relate the long-term volatilities to the macroeconomic environment. Thereafter, the DCC-MIDAS specifications that focus on the long-run correlations are discussed.

5.1 Determinants of long-term volatilities

Tables 2 and 3 present the estimates for the various stock and oil GARCH-MIDAS models. In addition to the models which include the macroeconomic variables, we consider the stationary GARCH(1,1) with constant unconditional variance as our benchmark specification. Since the serial correlation in daily stock and oil returns is negligible, we choose $\mu_{i,t} = \mu_i$ in both conditional means. The optimal lag length is found to be $K_v = 36$ for

⁵The NAI is a standardized weighted average of 85 monthly indicators of national economic activity including figures that represent (i) production and income, (ii) employment, unemployment, and hours, (iii) personal consumption and housing, and (iv) sales, orders and inventories. The NAI is computed and published by the Federal Reserve Bank of Chicago. Positive realizations indicate growth above trend, while negative realizations indicate growth below trend. The variables IP, NFP, and UR are among the indicators used for the computation of the NAI.

⁶The LI predicts the six-month growth rate of the US coincident index based on variables that lead the economy including housing permits, unemployment insurance claims, delivery times from the ISM manufacturing survey, and the term spread. The LI is published by the Federal Reserve Bank of Philadelphia.

both markets, i.e. our specifications cover three MIDAS lag years. However, all results are robust to moderate changes in K_v . We compare the fit of the different models by means of the Akaike and Bayesian information criteria (AIC and BIC).⁷ As an alternative measure of the fit of the various specifications, we evaluate the quality of volatility forecasts by means of a Mincer-Zarnowitz (MZ) regression. That is, for each model specification we regress the daily squared returns on a constant and the corresponding conditional variance forecast:

$$r_{i,t}^2 = c_i + \phi_i \cdot \hat{h}_{i,t} + \xi_{i,t}, \quad i = 1, 2. \quad (12)$$

We report the R^2 of the MZ regression and the p -value of the F -statistic for the joint hypothesis $c_i = 0$ and $\phi_i = 1$.

The constant μ_i is significant in all stock return models, but insignificant in the oil return specifications. In all cases the estimated α_i and β_i parameters are highly significant. Interestingly, while the α_i (β_i) parameters are estimated to be slightly higher (lower) in the stock than in the oil market, the sum $\alpha_i + \beta_i$ is almost identical in both markets and always less than one. That is, in all specifications the short-run volatility component is mean-reverting to the long-run trend. Next, we discuss the estimated θ_i and ω_i parameters individually for the two markets.

Tables 2 and 3 about here

Since the macroeconomic determinants of long-term stock market volatility have been investigated in Engle et al. (2012) and Conrad and Loch (2012) already, we only briefly summarize our findings which are very much in line with theirs. Table 2 shows that each macroeconomic variable has a significant effect on long-term stock market volatility. For IP, NFP, NAI, and LI the estimated coefficient $\hat{\theta}_1$ is negative and highly significant, while it is positive and highly significant in case of UR. Since the sign of θ_1 measures whether an increase of the respective variable leads to an upswing or downswing in long-run volatility, the estimates imply that higher (lower) levels of economic activity lead to a reduction (rise) in long-term stock market volatility. According to the R^2 's from the MZ regressions, the model based on IP appears to perform best. In addition, all GARCH-MIDAS-X models except the UR specification yield slightly higher R^2 's than the benchmark GARCH(1,1). Similarly, all GARCH-MIDAS-X models are preferred over

⁷Note that all GARCH-MIDAS-X models include the same number of parameters and, hence, the AIC and BIC will lead to the same ranking. However, the benchmark GARCH(1,1) model includes two parameters less.

the benchmark GARCH(1,1) by the AIC, but not by the BIC. In short, our results reconfirm the observation in Engle et al. (2012) and Conrad and Loch (2012) that long-term stock market volatility behaves counter cyclically and that the GARCH-MIDAS-X models outperform the simple GARCH(1,1).

In Table 3 we turn to the analysis of the macroeconomic determinants of the long-term oil return volatility. The estimates of θ_2 suggest that long-term oil return volatility is closely linked to each of the macroeconomic variables describing the current stance of the economy as well as the future economic outlook. In particular, the results imply that downturns in U.S. economic activity, i.e. decreases in IP, NFP, NAI, and LI and increases in UR lead to higher levels of long-term oil return volatility. While Kilian (2008, 2009), Kilian and Murphy (2013) and Alquist et al. (2013) have provided ample evidence for the notion that changes in economic activity predict oil prices, our finding that U.S. economic activity also precedes changes in long-term oil return volatility adds a new insight. Given the positive relation between aggregate demand shocks and the level of the oil price which was established in the previous literature, our finding of a counter cyclical behavior of long-term oil return volatility is very much in line with the observation in stock markets that volatility is low during phases of increasing prices but high during phases of decreasing prices. That is, good news on the macroeconomy is also good news for the oil market, i.e. increases the oil price and at the same time reduces oil return volatility.

Lastly, all GARCH-MIDAS-X models achieve a better fit than the GARCH(1,1), both in terms of the AIC and the BIC. The best model according to the information criteria is the one based on the LI. This model also leads to the highest R^2 in the MZ regressions. While we cannot reject the null hypothesis of unbiased forecasts for all GARCH-MIDAS-X models, the hypothesis is clearly rejected (at the 5% level) in case of the benchmark GARCH(1,1).

Figure 2 shows the GARCH-MIDAS-LI estimates of the annualized monthly long-term volatility components for the two markets. While the level of oil return volatility is about twice as high as the one of the stock returns, the evolution of the two components is very similar across markets. The observation that the macroeconomic environment affects long-term oil and stock volatility in almost the same way is very interesting. Our finding suggests that the long-term second moment of oil price returns shares a common component with that of stock returns which reflects the state of the U.S. business cycle.

Figure 2 about here

5.2 Determinants of long-term correlations

Next, we analyze the macroeconomic determinants of the long-term oil-stock correlation. We consider two benchmark specifications. The first natural benchmark is the DCC-GARCH model. The second benchmark is the Colacito et al. (2011) specification, which uses backward-looking monthly realized correlations as explanatory variables. For both the DCC and the DCC-MIDAS-RC we employ the standardized residuals from the simple GARCH(1,1) model. In the general DCC-MIDAS-X specifications we replace the realized correlations with key macroeconomic figures. For these models the volatility-adjusted residuals are obtained from the corresponding GARCH-MIDAS-X models. As in the case of the long-term volatilities, we find that the optimal lag length is equal to three MIDAS lag years, i.e. we choose $K_c = 36$. In addition to the AIC and BIC, we use the R^2 of the following MZ regression to evaluate the various model specifications in predicting covariances:

$$r_{1,t}r_{2,t} = c_{12} + \phi_{12}\hat{h}_{12,t} + \xi_{12,t}, \quad (13)$$

where $\hat{h}_{12,t} = \hat{\rho}_{12,t}\sqrt{\hat{h}_{1,t}\hat{h}_{2,t}}$. Again, we test the joint hypothesis $c_{12} = 0$ and $\phi_{12} = 1$.

Table 4 presents the estimation results. Clearly, in all specifications the estimated parameters a and b are highly significant and sum up to a value of less than one. That is, the quasi-correlations are mean-reverting either to the unconditional correlation in the DCC-GARCH case or to the long-term correlation in the various DCC-MIDAS-X specifications. The estimates of θ_{12} indicate that all macroeconomic variables significantly affect the long-run oil-stock correlation. In line with our analysis in Section 5.1, we find negative θ_{12} coefficients on IP, NFP, NAI, and LI, while the coefficient on UR is positive. The estimates imply that a contraction of macroeconomic activity leads to an increase of the long-term correlation.

Table 4 about here

According to the AIC and BIC, all DCC-MIDAS-X specifications are superior relative to the DCC-GARCH. Hence, there is convincing evidence in favor of the component models, which allow for time-varying long-term correlations. While the DCC-MIDAS-UR model achieves the highest R^2 in the MZ regressions, the information criteria favor the model based on the LI. The null hypothesis of unbiased covariance forecasts cannot be rejected for any specification. The fact the DCC-MIDAS-X models are also preferred (both in terms of information criteria and MZ R^2) to the DCC-MIDAS-RC suggests

that the various macroeconomic variables carry information on the evolution of the long-term correlation beyond that included in past realized correlations. Next, we explain how the forward looking properties of the macroeconomic variables which gauge future economic activity as well as inflationary pressures (and thereby future monetary policy) are particularly relevant for anticipating changes in the oil-stock correlation.

Figure 3 shows the estimated dynamics of the short- and long-run correlations based on the DCC-MIDAS-LI specification together with a rolling-window of yearly realized correlations. First, although the unconditional correlation between stock and oil returns was found to be 0.14, the figure shows that there is substantial time-variation in the realized correlations with prolonged periods of positive or negative correlations. While the short-run correlation closely follows the behavior of the realized correlations, the long-run correlation evolves much more smoothly. Both the realized correlations as well as the short-run correlations follow this long-run trend component.

Figure 3 about here

To provide an economic interpretation of the cyclical pattern in the evolution of the correlation dynamics we refer to Figure 4, which depicts the long-term component along with the LI. First, the figure clearly shows an inverse relationship between the LI and the long-term oil-stock correlation, which was already evident from the negative θ_{12} estimate in Table 4. That is, the oil-stock correlation is increasing (decreasing) when the LI is declining (rising). Interestingly, when the LI decreases and turns negative before and during the 2001 and 2007-2009 recessions the long-term correlation steeply increases, while it decreases more gradually in the aftermath of the recessions. On the other hand, the long period of strong growth from 1994 to 1999 is accompanied by a period of negative oil-stock correlations. Our empirical evidence for a counter cyclical oil-stock correlation is perfectly in line with the recent evidence in Kilian (2009) in favor of a positive oil-growth relation. Kilian and Park (2009) argue that in an early phase of an expansion increasing oil prices may not have negative effects on the stock market. This is because in the short-run the positive effect of higher economic activity on expected future cash flows dominates and, hence, the oil-stock correlation will be positive. However, in the long-run the negative effect of increasing oil prices on corporate cash flows will dominate and, therefore, the oil-stock correlation will decrease or even turn negative.

Figure 4 about here

The long-term correlation in Figure 4 very much supports these views. Before and during both recessions bad news on the LI is associated with sharply decreasing stock and oil prices and, therefore, a positive oil-stock correlation. The fact that the correlation turns positive and increases well before both recessions is remarkable and suggests that the long-term oil-stock correlation may itself be used as an early recession indicator. During the recovery phases in 2002-2003 and 2010-2011 the improvement in the LI leads to increasing oil prices and, at the same time, to upward revisions concerning firms' expected dividends and cash flows. In these periods the oil-stock correlation remains positive, but smoothly decreases. The same rationale also applies to the first year of our sample, which falls into the recovery period after the recession of 1990/91. Finally, during the years 1994-1999 and 2004-2006 the LI signals strong growth for a protracted period, which again should positively affect oil prices. However, the (expected) oil price increases now dampen the outlook for future corporate cash flows, i.e. during these periods the good news on the macroeconomy – through the indirect effect via increasing oil prices – turns into bad news for the stock market. Alternatively, the negative effect might also work via interest rates. When the economy is already close to full employment, good news on the LI could signal higher future interest rates and, hence, be bad news for the stock market. During these strong boom phases the negative effect dominates and leads to a decreasing or negative long-run oil-stock correlation.

Since the evolution of the long-term correlation is purely driven by variables which represent U.S. aggregate demand, deviations of the short-term component from the long-run trend must be related to other factors which affect stock and/or oil returns. Typical oil related factors would be oil supply shocks or oil-market specific demand shocks such as precautionary demand or speculative demand shocks. Specifically, the temporary deviation in 2002/03 (see Figure 3) can be explained as a combination of the Venezuelan oil supply crisis and precautionary demand provoked by the second Iraq war (see Kilian and Murphy, 2013). Similarly, the drop in the short-term component in 2011/02-2011/04 can be related to the Libyan crisis and political turmoil in North Africa.⁸ Another example would be the positive correlation signaled by the short-term component as well as the realized correlations around 1998/99. Following the Asian and Russian financial crises, this positive short-term correlation can be explained by a simultaneous decline in oil and stock prices. Nevertheless, the fact that these deviations occur only for relatively short

⁸On February 22nd 2011, for instance, oil returns spiked up by 8%, whereas stock market returns went down by 2%.

periods suggests that the oil-stock correlation can be largely explained by U.S. economic activity for most of the time.

A particularly important conclusion that can be drawn from the time-varying oil-stock correlation is that regressions of stock returns on oil price changes are likely to be misleading, since the result will depend on the state of the economy. This insight may explain the controversial empirical findings on the oil-stock relationship and agrees with the arguments put forward in Kilian and Park (2009).

Next, we discuss the MIDAS lag structure and its implications more closely. Recall that the higher ω_{12} the more weight will be given to the more recent observations of the macro variable and, hence, the faster the weights will decline to zero. Table 4 reveals that the lowest ω_{12} is estimated for IP and the highest for NFP. Since the DCC-MIDAS-LI model produced the best fit for the correlations, in Figure 5 we plot the corresponding weighting function. For comparison, we also display the weighting functions for the GARCH-MIDAS-LI models for the stock and oil market. The figure shows that the weighting function of the correlation model is nearly linear while the weighting functions of the volatility specifications are rapidly declining.⁹ This in turn implies that changes in the LI have a much more persistent effect on the long-run correlation than on the long-run volatilities.¹⁰

Figure 5 about here

In the previous considerations we mainly focused on the DCC-MIDAS-LI specification to explain the dynamic behavior of the slowly-moving long-run correlation component. However, Table 4 clearly reveals that the fit of the DCC-MIDAS-X specifications with IP, NFP, UR, and NAI are only slightly inferior. Figure 6 displays the estimated long-run correlations from the corresponding specifications. The figure illustrates nicely that the long-term components of all specifications follow the same pattern and, hence, further support our argument that the long-term oil-stock correlation is counter cyclical. Note that the exceptional deviation in the long-term correlation component predicted by IP for October 2005 can be traced back to a significant contraction in industrial production one

⁹We additionally estimated the models including a weighting scheme with two parameters, hereby relaxing the assumption of strictly decreasing weights. However, including an unrestricted weighting scheme did not lead to significant improvements in the value of the maximized log likelihood and the resulting weighting schemes were still strictly decreasing.

¹⁰Similar results are obtained for the other macroeconomic variables but omitted for reasons of brevity.

month earlier. This is not reflected to such a strong extent in the other macroeconomic figures (compare Figure 1).

Figure 6 about here

6 Model evaluation and hedging performance

Although the main focus of our analysis lies on the macroeconomic determinants of the long-term oil price return volatility as well as the long-term oil-stock correlation, our findings might also have important implications for portfolio choice, hedging decisions or risk management. While we have already employed the R^2 's of the different MZ regressions for model ranking, we now have a closer look at the forecasting performance of the different models for the entire conditional covariance matrix \mathbf{H}_t . Since a full-fledged out-of-sample analysis is beyond the scope of the current paper, we simply focus on in-sample results. Following Laurent et al. (2012, 2013) we apply two robust loss functions, i.e. loss functions that deliver the same ordering whether the evaluation is based on the true conditional covariance matrix or an unbiased proxy of it. The first loss function is the Euclidean distance which equally weights the variances and covariances:

$$L_t^E = (r_{1,t}^2 - \hat{h}_{1,t})^2 + (r_{2,t}^2 - \hat{h}_{2,t})^2 + (r_{1,t}r_{2,t} - \hat{h}_{12,t})^2$$

The second one is based on the Frobenius distance and double counts the loss associated with the conditional covariance:

$$L_t^F = (r_{1,t}^2 - \hat{h}_{1,t})^2 + (r_{2,t}^2 - \hat{h}_{2,t})^2 + 2(r_{1,t}r_{2,t} - \hat{h}_{12,t})^2$$

In Table 5, we report for each model the average value of the two loss functions. In addition, for each DCC-MIDAS-X model we test whether the average loss is significantly different from the average loss of the DCC benchmark model. Panel A presents results for the full sample, while Panel B covers the subsample of the financial crisis in the years 2007-2009. In case of a positive difference, forecasts from the DCC-MIDAS-X model are superior to those from the benchmark model.

For the full sample, the differences in both loss functions are significant for all DCC-MIDAS-X models except the one based on IP. To the contrary, the DCC-MIDAS-RC model does not lead to a significant improvement over the simple DCC. Unsurprisingly, during the financial crisis period the average losses more than double in comparison to the

full sample. During this period we only find a significant improvement over the DCC for the model based on the LI. This somewhat disappointing outcome may be due to the fact that during the crisis the forecast quality of all models deteriorated dramatically and it became increasingly difficult to distinguish between them. Another potential explanation could be that during the crisis the quality of our proxies, i.e. the squared returns and the product of the oil and stock returns, for the true conditional volatilities and covariances has declined.

As an alternative approach to evaluate the forecast performance, we consider the problem of hedging a long position of one dollar in the stock market by a short position of $\beta_{12,t}$ dollars in the oil market. The optimal hedge portfolio is given by (see Kroner and Sultan, 1993):

$$r_t^{PF} = r_{1,t} - \beta_{12,t} \cdot r_{2,t}, \quad \text{with} \quad \beta_{12,t} = \frac{\hat{h}_{12,t}}{\hat{h}_{2,t}}.$$

We then compare the average portfolio variance based on the volatility and covariance forecasts from the DCC-MIDAS-X models with those from the DCC model. The results in Table 5 suggest that the DCC-MIDAS-X models lead to significantly lower portfolio variances compared to the DCC in both the full sample as well as the crisis subsample. Although, the forecasting results are very promising for potential financial applications, a natural avenue for future research would be to confirm our in-sample findings in a more detailed out-of-sample analysis.

7 Conclusions

We investigate the effect of changes in the U.S. macroeconomic environment on the long-term volatilities and correlations in crude oil and U.S. stock price returns. First, our results show that the long-term volatilities in both markets share a common component that reflects the state of the U.S. business cycle. Second, we extend the two-component DCC-MIDAS model of Colacito et al. (2011) by allowing the slowly-moving long-term correlation component to be determined endogenously by the variation of key macroeconomic figures. We show that changes in macroeconomic variables, which reflect the current stance of the economy as well as the future economic outlook, can anticipate counter cyclical fluctuations in the long-term correlation. More specifically, our model predicts a negative correlation during prolonged periods of strong economic expansions, while a positive correlation is observed during recessions and recoveries.

Our results provide further evidence for the argument put forward in Barsky and Kilian (2002, 2004) and Kilian (2008, 2009), among others, that oil price changes should not be considered exogenous with respect to U.S. and global macroeconomic conditions. However, while previous studies focused on a relationship in levels, our analysis shows that there is also feedback from the level of the macro variables to the second moment of the oil price. In addition, our MIDAS approach allows us to establish a link between low frequency data on U.S. economic activity and high frequency oil-stock return correlations, whereas previous evidence in Kilian and Park (2009) was based low frequency data.

References

- [1] Alquist, R., Kilian, L., and Vigfusson, R. J., 2013. Forecasting the price of oil. Forthcoming in: Elliott, G., Timmermann, A. (Eds.). *Handbook of Economic Forecasting*, 2, Amsterdam: North-Holland.
- [2] Barsky, R. B., and Kilian, L., 2002. Do we really know that oil caused the Great Stagflation? A monetary alternative. In: Bernanke, B. S., Rogoff, K. (Eds.). *NBER Macroeconomics Annual 2001*, MIT Press, Cambridge, MA, 137-183.
- [3] Barsky, R., and Kilian, L., 2004. Oil and the macroeconomy since the 1970s. *Journal of Economic Perspectives* 18, 115-134.
- [4] Bollerslev, T., and Wooldridge, J., 1992. Quasi-maximum likelihood estimation and inference in dynamic models with time-varying covariances. *Econometric Reviews* 11, 143-172.
- [5] Christodoulakis, G. and Satchell, S., 2002. Correlated ARCH (CorrARCH): Modelling the time-varying conditional correlation between financial asset returns. *European Journal of Operational Research* 139, 351-370.
- [6] Colacito, R., Engle, R., and Ghysels, E., 2011. A component model for dynamic correlations. *Journal of Econometrics* 164, 45-59.
- [7] Conrad, C., 2010. Non-negativity conditions for the hyperbolic GARCH model. *Journal of Econometrics* 157, 441-457.
- [8] Conrad, C., and Loch, K., 2012. Anticipating long-term stock market volatility. Working paper. Heidelberg University.
- [9] Davidson, J., 2004. Moment and memory properties of linear conditional heteroscedasticity models, and a new model. *Journal of Business and Economic Statistics* 22, 16-19.
- [10] Ding, Z., and Granger, C., 1996. Modeling volatility persistence of speculative returns: A new approach. *Journal of Econometrics* 73, 185-215.
- [11] Engle, R., 2002. Dynamic conditional correlation - a simple class of multivariate GARCH models. *Journal of Business and Economic Statistics* 20, 339-350.

- [12] Engle, R., Ghysels, E., and Sohn, B., 2012. Stock market volatility and macroeconomic fundamentals. *Review of Economics and Statistics*, forthcoming.
- [13] Engle, R., and Lee, G., 1999. A permanent and transitory component model of stock return volatility. In: Engle, R., White, H. (Eds.). *Cointegration, Causality and Forecasting: A Festschrift in Honor of Clive W.J. Granger*. Oxford University Press.
- [14] Engle, R., Rangel, J., 2008. The spline GARCH model for unconditional volatility and its global macroeconomic causes. *Review of Financial Studies* 21, 1187-1222.
- [15] Filis, G., Degiannakis, S., Floros, C., 2011. Dynamic correlation between stock market and oil prices: The case of oil-importing and oil-exporting countries. *International Review of Financial Analysis* 20, 152-164.
- [16] Ghysels, E., Santa-Clara, P., and Valkanov, R., 2005. There is a risk-return trade-off after all. *Journal of Financial Economics* 76, 509-548.
- [17] Ghysels, E., Sinko, A., and Valkanov, R., 2007. MIDAS regressions: Further results and new directions. *Econometric Reviews* 26, 53-90.
- [18] Hamilton, J., 1983. Oil and the macroeconomy since World War II. *Journal of Political Economy* 91, 228-248.
- [19] Hamilton, J., 1985. Historical causes of postwar oil shocks and recessions. *Energy Journal* 6, 97-116.
- [20] Hamilton, J., 2003. What is an oil shock? *Journal of Econometrics* 113, 363-398.
- [21] Hamilton, J., 2008. Oil and the macroeconomy. In: Durlauf, S., and Blume, L. (Eds.). *New Palgrave Dictionary of Economics*. 2nd edition. Palgrave MacMillan Ltd.
- [22] Huang, R., Masulis, R., and Stoll, H., 1996. Energy shocks and financial markets. *Journal of Futures Markets* 16, 1-27.
- [23] Jones, C., and Kaul, G., 1996. Oil and the stock markets. *Journal of Finance* 51, 463-491.
- [24] Kilian, L., 2008. The economic effects of energy price shocks. *Journal of Economic Literature* 46, 871-909.

- [25] Kilian, L., 2009. Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review* 99, 1053-1069.
- [26] Kilian, L., and Murphy, D.P., 2013. The role of inventories and speculative trading in the global market for crude oil. *Journal of Applied Econometrics*, forthcoming.
- [27] Kilian, L., and Park, C., 2009. The impact of oil price shocks on the U.S. stock market. *International Economic Review* 50, 1267-1287.
- [28] Kilian, L., and Vega, C., 2011. Do energy prices respond to U.S. macroeconomic news? A test of the hypothesis of predetermined energy prices. *Review of Economics and Statistics* 93, 660-671.
- [29] Kroner, K. F., and Sultan, J., 1993. Time-varying distributions and dynamic hedging with foreign currency futures. *Journal of Financial and Quantitative Analysis* 28, 535-551.
- [30] Laurent, S., Rombouts, J. V. K., and Violante, F., 2012. On the forecasting accuracy of multivariate GARCH models. *Journal of Applied Econometrics* 27, 934-955.
- [31] Laurent, S., Rombouts, J. V. K., and Violante, F., 2013. On loss functions and ranking forecasting performances of multivariate volatility models. *Journal of Econometrics* 173, 1-10.
- [32] Miller, J., and Ratti, R., 2009. Crude oil and stock markets: Stability, instability, and bubbles. *Energy Economics* 31, 559-568.
- [33] Nandha, M., and Faff, R., 2008. Does oil move equity prices? A global view. *Energy Economics* 30, 986-997.
- [34] Schwert, W., 1989. Why does stock market volatility change over time? *Journal of Finance* 44, 1115-1153.
- [35] Wei, C., 2003. Energy, the stock market, and the Putty-Clay investment model. *American Economic Review* 93, 311-323.

8 Appendix

8.1 Tables

Table 1: Descriptive Statistics

Variable	Obs	Min	Max	Mean	Std. Dev.*	Skewness	Kurtosis
Panel A: Daily return data (Jan 4, 1993 - Nov 30, 2011)							
Oil (WTI)	4743	-17.09	16.41	0.03	39.21	-0.19	7.73
CRSP	4743	-8.96	11.35	0.04	19.53	-0.11	10.66
Panel B: Monthly macro data (Jan 1990 - Nov 2011)							
<i>Current stance of the economy</i>							
IP	263	-4.30	2.10	0.16	0.67	-1.72	11.52
NFP	263	-0.62	0.41	0.07	0.18	-1.16	5.17
UR	263	-0.50	0.50	0.01	0.16	0.39	3.88
<i>Future economic outlook</i>							
NAI	263	-4.55	1.52	-0.17	0.86	-1.82	8.48
LI	263	-3.03	2.42	0.99	0.98	-1.67	6.69
Notes: *The standard deviations are annualized for the daily return series.							

Table 2: GARCH-MIDAS-X models: stock market

Variable	μ_1	α_1	β_1	m_1	θ_1	ω_1	LLF	BIC	AIC	R^2
<i><u>Current stance of the economy</u></i>										
IP	0.0672*** (0.0121)	0.0827*** (0.0121)	0.9070*** (0.0135)	0.4005* (0.2283)	-0.9588** (0.4245)	3.2001* (1.8916)	-6589.33	2.7893	2.7811	22.3 [0.645]
NFP	0.0677*** (0.0120)	0.0863*** (0.0128)	0.9003*** (0.0147)	0.4052** (0.1989)	-2.3472*** (0.5559)	10.1923 (6.7740)	-6586.41	2.7880	2.7799	21.8 [0.557]
UR	0.0675*** (0.0120)	0.0838*** (0.0119)	0.9034*** (0.0136)	0.1642 (0.1954)	4.2234*** (0.9428)	5.5672* (3.0725)	-6585.59	2.7877	2.7795	21.6 [0.546]
<i><u>Future economic outlook</u></i>										
NAI	0.0674*** (0.0121)	0.0848*** (0.0124)	0.9028*** (0.0144)	0.1374 (0.2073)	-0.5545*** (0.1401)	6.1885 (4.9129)	-6587.38	2.7884	2.7803	22.0 [0.603]
LI	0.0672*** (0.0120)	0.0851*** (0.0125)	0.9011*** (0.0146)	0.7015*** (0.2157)	-0.4850*** (0.1165)	6.1449 (6.0315)	-6585.34	2.7876	2.7794	22.0 [0.593]
<i><u>Benchmark model</u></i>										
GARCH(1,1)	0.0660*** (0.0121)	0.0810*** (0.0118)	0.9109*** (0.0129)	0.2467 (0.2533)	-	-	-6592.08	2.7868	2.7814	21.7 [0.547]

Notes: The numbers in parentheses are Bollerslev-Wooldridge (2002) robust standard errors. ***, **, * indicate significance at the 1 %, 5 %, and 10 % level. LLF is the value of the maximized likelihood function, BIC is the Bayesian information criterion and AIC is the Akaike information criterion.

We run Mincer-Zarnowitz regressions of the form

$$r_{1,t}^2 = c_1 + \phi_1 \cdot \hat{h}_{1,t} + \xi_{1,t},$$

where daily squared stock market returns are regressed on a constant and the conditional variance estimate obtained from the GARCH-MIDAS-X or the benchmark GARCH(1,1) model. For each regression we report the R^2 percentage value and in brackets the p -value from the F -statistic of a Wald Test with the null hypothesis $c_1 = 0$ and $\phi_1 = 1$. The bold numbers indicate the models with the lowest (highest) values of the information criteria (MZ R^2).

Table 3: GARCH-MIDAS-X models: oil market

Variable	μ_2	α_2	β_2	m_2	θ_2	ω_2	LLF	BIC	AIC	R^2
<i>Current stance of the economy</i>										
IP	0.0491 (0.0331)	0.0583*** (0.0160)	0.9224*** (0.0227)	1.8464*** (0.1360)	-0.5121** (0.2135)	7.5161*** (2.3741)	-10586.33	4.4747	4.4665	9.8 [0.576]
NFP	0.0477 (0.0331)	0.0601*** (0.0159)	0.9214*** (0.0213)	1.8629*** (0.1473)	-1.0675** (0.4418)	15.9548*** (5.0328)	-10587.08	4.4750	4.4668	10.3 [0.922]
UR	0.0489 (0.0327)	0.0575*** (0.0145)	0.9232*** (0.0202)	1.7372*** (0.1268)	2.0857*** (0.6705)	12.6712*** (4.2630)	-10584.25	4.4738	4.4656	10.3 [0.938]
<i>Future economic outlook</i>										
NAI	0.0486 (0.0331)	0.0574*** (0.0158)	0.9244*** (0.0218)	1.7227*** (0.1364)	-0.2952*** (0.1105)	14.9278*** (5.2959)	-10585.90	4.4745	4.4663	10.1 [0.684]
LI	0.0465 (0.0330)	0.0558*** (0.0161)	0.9253*** (0.0226)	2.0755*** (0.1632)	-0.3111*** (0.0915)	21.5657** (8.9818)	-10582.68	4.4732	4.4650	10.6 [0.875]
<i>Benchmark model</i>										
GARCH(1,1)	0.0363 (0.0394)	0.0518*** (0.0155)	0.9523*** (0.0124)	-0.0074 (0.5920)	-	-	-10620.82	4.4857	4.4802	10.0 [0.029]

Notes: We run Mincer-Zarnowitz regressions of the form

$$r_{2,t}^2 = c_2 + \phi_2 \cdot \hat{h}_{2,t} + \xi_{2,t},$$

where daily squared oil market returns are regressed on a constant and the conditional variance estimate obtained from the GARCH-MIDAS-X or the benchmark GARCH(1,1) model. For each regression we report the R^2 percentage value and in brackets the p -values from the F -statistic of a Wald Test with the null hypothesis $c_2 = 0$ and $\phi_2 = 1$. Otherwise see Notes of Table 2.

Table 4: DCC-MIDAS-X models: stock and oil market

Variable	a	b	m_{12}	θ_{12}	ω_{12}	LLF	BIC	AIC	R^2
<i>Current stance of the economy</i>									
IP	0.0188*** (0.0060)	0.9712*** (0.0102)	0.2038*** (0.0582)	-0.6931*** (0.1885)	1.7758* (0.9689)	-4678.15	1.9816	1.9748	6.0 [0.598]
NFP	0.0190*** (0.0063)	0.9708*** (0.0111)	0.2136*** (0.0571)	-1.4176*** (0.3316)	4.0474 (2.7107)	-4670.72	1.9784	1.9716	6.1 [0.755]
UR	0.0214*** (0.0056)	0.9620*** (0.0103)	0.0505 (0.0360)	3.0976*** (0.6679)	1.9275** (0.8945)	-4668.60	1.9776	1.9707	6.6 [0.772]
<i>Future economic outlook</i>									
NAI	0.0190*** (0.0055)	0.9661*** (0.0102)	0.0361 (0.0369)	-0.3647*** (0.0754)	2.1842* (1.1314)	-4669.00	1.9777	1.9709	6.4 [0.831]
LI	0.0189*** (0.0060)	0.9702*** (0.0103)	0.3615*** (0.0810)	-0.2647*** (0.0628)	2.3472 (1.5240)	-4663.08	1.9752	1.9684	6.5 [0.865]
<i>Benchmark models</i>									
DCC-RC	0.0236*** (0.0055)	0.9565*** (0.0102)	0.0314 (0.0348)	0.7944** (0.3161)	5.6739* (3.3987)	-4749.66	2.0117	2.0049	5.2 [0.283]
DCC	0.0203*** (0.0067)	0.9751*** (0.0095)	-	-	-	-4763.60	2.0123	2.0095	5.1 [0.200]

Notes: We run Mincer-Zarnowitz regressions of the form

$$r_{1,t} \cdot r_{2,t} = c_{12} + \phi_{12} \cdot \hat{\rho}_{12,t} \sqrt{\hat{h}_{1,t} \cdot \hat{h}_{2,t}} + \xi_{12,t},$$

where the product of daily stock and oil market returns is regressed on a constant and the conditional covariance estimate from the DCC-MIDAS-X model. For each regression we report the R^2 percentage value and in brackets the p -value from the F -statistic of a Wald Test with the null hypothesis $c_{12} = 0$ and $\phi_{12} = 1$. Otherwise see Notes of Table 2.

Table 5: Model evaluation

Variable	Euclidean distance		Frobenius distance		Hedge portfolio	
	loss	difference	loss	difference	variance	difference
Panel A: Full sample (Jan 1993 - Nov 2011)						
<u>Current stance of the economy</u>						
IP	264.604	1.907 (1.158)	287.063	2.136 (1.311)	1.399	0.016* (1.896)
NFP	263.328	3.183** (2.03)	285.744	3.455** (2.181)	1.398	0.017* (1.892)
UR	263.078	3.433** (2.152)	285.394	3.805** (2.344)	1.394	0.021* (1.876)
<u>Future economic outlook</u>						
NAI	263.738	2.772** (2.008)	286.073	3.126** (2.223)	1.395	0.020* (1.919)
LI	262.477	4.034*** (2.895)	284.801	4.398*** (3.054)	1.396	0.019** (1.982)
<u>Benchmark</u>						
DCC-RC	266.505	0.006 (0.153)	289.187	0.012 (0.153)	1.414	0.001 (0.538)
DCC	266.511	-	289.199	-	1.415	-
Panel B: Financial crisis (Jan 2007 - Dec 2009)						
<u>Current stance of the economy</u>						
IP	629.624	0.221 (0.023)	720.357	1.628 (0.175)	3.012	0.100** (2.151)
NFP	621.292	8.553 (0.967)	711.732	10.252 (1.153)	3.004	0.109** (2.193)
UR	622.677	7.168 (0.824)	712.467	9.517 (1.080)	2.979	0.134** (2.196)
<u>Future economic outlook</u>						
NAI	624.536	5.309 (0.715)	714.491	7.493 (0.995)	2.987	0.126** (2.153)
LI	618.037	11.808 (1.592)	707.883	14.101** (1.846)	2.997	0.116** (2.240)
<u>Benchmark</u>						
DCC-RC	629.867	-0.023 (-0.098)	722.030	-0.045 (-0.098)	3.117	-0.004 (-0.507)
DCC	629.845	-	721.984	-	3.113	-

Notes: For each DCC-MIDAS model we report the average of the Euclidean and Frobenius loss functions:

$$L_t^E = (r_{1,t}^2 - \hat{h}_{1,t})^2 + (r_{2,t}^2 - \hat{h}_{2,t})^2 + (r_{1,t}r_{2,t} - \hat{h}_{12,t})^2,$$

$$L_t^F = (r_{1,t}^2 - \hat{h}_{1,t})^2 + (r_{2,t}^2 - \hat{h}_{2,t})^2 + 2(r_{1,t}r_{2,t} - \hat{h}_{12,t})^2,$$

and the average difference relative to the benchmark DCC model along with values of the corresponding t -statistic. For each DCC-MIDAS model we calculate the optimal hedge portfolio

$$r_t^{PF} = r_{1,t} - \beta_{12,t} \cdot r_{2,t}, \quad \text{with} \quad \beta_{12,t} = \frac{\hat{h}_{12,t}}{\hat{h}_{2,t}},$$

and report its average variance. The average variance for the portfolio consisting only of stock returns amounts to 1.507 for the full sample and to 3.438 for the subsample. We calculate the average difference of each variance relative to the DCC model and the corresponding t -statistic. ***, **, * indicate significance at the 1 %, 5 %, and 10 % level.

8.2 Figures

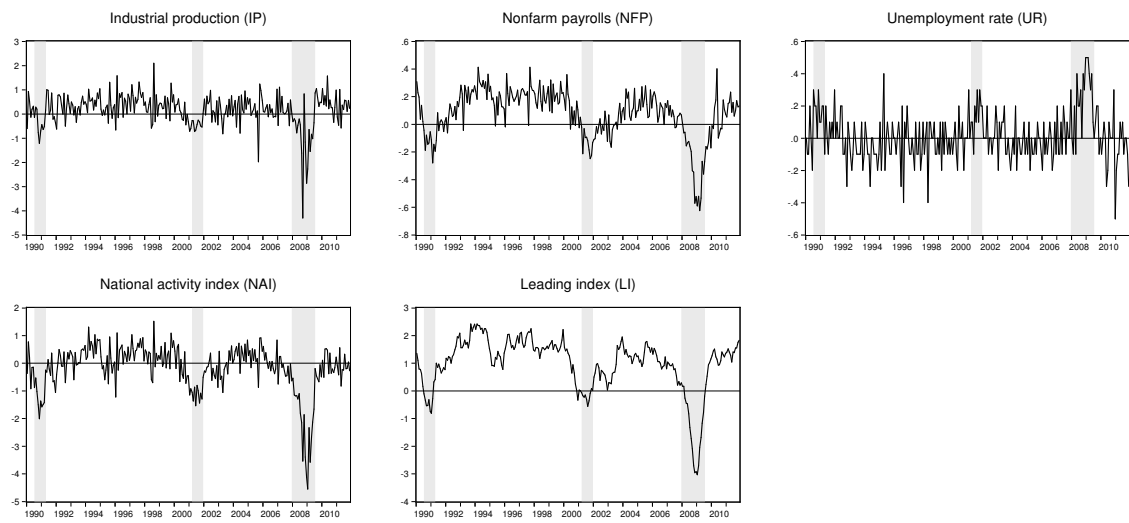


Figure 1: The figure shows the development of the macroeconomic explanatory variables. Shaded areas represent NBER recession periods.




Figure 2: The figure shows the annualized monthly long-term volatility components (standard deviations) obtained from the GARCH-MIDAS-LI specification. The bold line refers to the stock market, the dashed line to the oil market. Shaded areas represent NBER recession periods.




Figure 3: The figure shows the DCC-MIDAS-LI estimates of the conditional oil-stock correlation (dashed line) and its long-term component (bold black line). The circles correspond to one-year rolling window realized correlations. Each series is shown at a monthly frequency, where monthly realizations are obtained by computing monthly averages. Shaded areas represent NBER recession periods.




Figure 4: The bold black line (left scale) represents the DCC-MIDAS-LI estimate of the long-term oil-stock correlation. The dashed line (right scale) corresponds to the LI. Shaded areas represent NBER recession periods.




Figure 5: The figure shows the estimated weighting functions for the long-term volatilities based on the GARCH-MIDAS-LI and for the long-term correlation based on the DCC-MIDAS-LI. While the bold black line refers to the long-term correlation, the light-gray and the dark-gray dashed lines refer to the long-term volatilities of CRSP and of oil price returns, respectively.




Figure 6: The figure shows the DCC-MIDAS-X estimates of the long-term oil-stock correlations for all macroeconomic variables. Shaded areas represent NBER recession periods.