Chapter 1
Wavelet–based correlation analysis of the key traded assets

Jozef Baruník and Evžen Kočenda and Lukas Vacha

Abstract This chapter reveals the time-frequency dynamics of the dependence among key traded assets – gold, oil, and stocks, in the long run, over a period of 26 years. Using both intra-day and daily data and employing a variety of methodologies, including a novel time-frequency approach combining wavelet-based correlation analysis with high-frequency data, we provide interesting insights into the dynamic behavior of the studied assets. We account for structural breaks and reveal a radical change in correlations after 2007-2008 in terms of time-frequency behavior. Our results confirm different levels of dependence at various investment horizons indicating heterogeneity in stock market participants' behavior, which has not been documented previously. While these key assets formerly had the potential to serve as items in a well-diversified portfolio, the events of 2007-2008 changed this situation dramatically.

Key words: time-frequency dynamics; gold; oil; stocks; high-frequency data; dynamic correlation; crisis; wavelets

---

Jozef Baruník
Theory and Automation, Academy of Sciences of the Czech Republic, Pod Vodarenskou Vezi 4, 182 00, Prague, Czech Republic
Institute of Economic Studies, Charles University, Opletalova 21, 110 00, Prague, Czech Republic, e-mail: barunik@utia.cas.cz

Evžen Kočenda
CERGE-EI, Charles University and the Czech Academy of Sciences, Politickych veznu 7, 11121 Prague, Czech Republic
CESifo, Munich; IOS Regensburg
The William Davidson Institute at the University of Michigan Business School
CEPR, London; and the Euro Area Business Cycle Network. e-mail: evzen.kocenda@cerge-ei.cz

Lukas Vacha (corresponding author)
Theory and Automation, Academy of Sciences of the Czech Republic, Pod Vodarenskou Vezi 4, 182 00, Prague, Czech Republic
Institute of Economic Studies, Charles University, Opletalova 21, 110 00, Prague, Czech Republic, e-mail: vachal@utia.cas.cz

1.1 Introduction, motivation, and related literature

In this chapter, we contribute to the literature by studying the dynamic relationship among gold, oil, and stocks in a time–frequency domain by employing a wavelet–based methodology. Considering the time-frequency domain offers new perspectives on the relationships among the assets and differentiates our contribution from much of the related literature. The time–frequency approach also enables us to uncover patterns underpinned by the investment potential derived from the ongoing financialization of commodities.

Traders in financial markets make their decisions over various horizons, for example, minutes, hours, days, or even longer periods such as months and years, as discussed by Ramsey (2002). Nevertheless, majority of the empirical literature studies the relationships in the time domain only aggregating the behavior across all investment horizons. Our analysis includes both time and time–frequency methods. The time domain tools we apply to measure correlations are the parametric DCC GARCH and nonparametric realized volatility. Although these two methods are fundamentally different, they both average the relationships over the full range of available frequencies and suffer from restricted application when analyzing non-stationary time series. In contrast, wavelets allow us to analyze time series within a time-frequency domain framework that allows for various forms of localization. Thus, when analyzing non-homogeneous and non-stationary time series, wavelet analysis is preferred because it is more flexible. For example, when considering stock markets, we can work with prices and thus study the dynamics of the dependencies at various investment horizons or frequencies at the same moment, where the lowest frequency will contain the trend component of the data. Therefore, we can determine short and long–term dependence structures. Wavelets are able to deliver valuable and unorthodox inferences in the fields of economics and finance, as evinced in recent applications, for example, by Faÿ et al (2009), Gallegati et al (2011), Vacha and Barunik (2012), Aguiar-Conraria et al (2012), or Graham et al (2013).

Our analysis is performed using data from a long period, 26 years, from 1987 to 2012. We conduct a thorough, wavelet-based analysis and uncover rich time-frequency dynamics in the relationships among gold, oil, and stocks. The selection of the three assets is motivated by the fact that gold and oil are the most actively traded commodities in the world. Similarly, to represent stocks, we use the S&P500, which is one of the most actively traded and comprehensive stock indices in the world. Gold is traditionally perceived as a store of wealth, especially with respect to periods of political and economic insecurity (Aggarwal and Lucey, 2007). However, gold is both a commodity and a monetary asset. Approximately 40% of newly mined gold is used for investment (Thompson Reuters GFSM, 2012). Unlike gold, oil is essential component of contemporary industrial economies, as reflected by the 88 million barrels consumed daily worldwide. As oil is a vital production input, its price is driven by distinct demand and supply shocks (Hamilton, 2009). Oil has also become financialized over time, as documented in Büyükşahin and Robe (2013). Fratzscher et al (2013) show that oil was not correlated with stocks until

2001, but as oil began to be used as a financial asset, the link between oil and other assets strengthened. Finally, stocks reflect the economic and financial development of firms and market perceptions of a companys standing; they also represent investment opportunities and a link to perceptions of aggregate economic development. Further, stock prices provide helpful information on financial stability and can serve as an indicator of crises (Gadanecz and Jayaram, 2009). Thus, a broad market index can be used to convey information on the status and stability of the economy. In our analysis, we consider the S&P 500, which is frequently used as a benchmark for the overall U.S. stock market. In our analysis, stocks complement the commodities of gold and oil to represent the financial assets traded by the modern financial industry.

What motivates our analysis of the links among the three assets above? The literature analyzing the dynamic correlations among assets proposes a number of important reasons that the issue should be investigated. An obvious motivation for analyzing co-movements is that substantial correlations among assets greatly reduce their potential to be included in a portfolio from the perspective of risk diversification. Even if assets in a portfolio initially exhibit low correlation, a potential change in correlation patterns represents an imperative to redesign such a portfolio. Both issues are also linked to the Modern Portfolio Theory (MPT) of Markowitz (1952). MPT assumes, among other things, that correlations between assets are constant over time. However, correlations between assets may well depend on the systemic relationships between them and change when these relationships change. Thus, evidence of time-varying correlations between assets substantially undermines MPT results and, more important, its use to protect investors from risk.

Empirical evidence on co-movements among assets may well depend on the choice of assets, technique employed, and the period under study. In a seminal study on co-movements in the monthly prices of unrelated commodities, Pindyck and Rotemberg (1990) find excess co-movement among seven major commodities, including gold and oil. However, the co-movements are measured in a rather simple manner as individual cross-correlations over the entire period (1960–1985). The excess co-movements were attributed to irrational or herding behavior in markets. Using a concordance measure, Cashin et al (1999) analyze the same set of commodities over the same period as Pindyck and Rotemberg (1990) and find no evidence for co-movements in the prices of the analyzed commodities. When they extend the period to 1957–1999, the co-movements are again absent and they contend that the entire notion of co-movements in the prices of unrelated commodities is a myth. A single exception is the co-movement in gold and oil prices that Cashin et al (1999) credit to inflation expectations and further provide evidence that booms in oil and gold prices often occur at the same time (Cashin et al, 2002). Still, it has to be noted that gold may well be traded independently from other assets on the pretext of being a store of value during downward market swings. Hence, it does not necessarily co-move with related or unrelated commodities.

An extension of the co-movement analysis to the time-frequency domain offers the potential for an interesting comparison of how investment horizons influence the diversification of market risk. The importance of various investment horizons for portfolio selection has been recognized by Samuelson (1989). In this respect, Marshall (1994) demonstrates that investor preferences for risk are inversely related to time and different investment horizons have direct implications for portfolio selection. Graham et al (2013) provide empirical evidence related to the issues studied in this chapter by studying the co-movements of various assets using wavelet coherence and demonstrating that at the long-term investment horizon co-movement among stocks and commodities increased at the onset of the 2007–2008 financial crisis. Thus, the diversification benefits of using these assets are rather limited.

With the above motivations and findings in mind, in this chapter we adopt a comprehensive approach and contribute to the literature by analyzing the prices of three assets that have unique economic and financial characteristics: the key commodities gold and oil and important stocks represented by the S&P 500 index. To this end, we consider a long period (1987–2012) at both intra-day and daily frequencies and an array of investment horizons to deliver a comprehensive study in the time-frequency domain based on wavelet analysis. Our key empirical results can be summarized as follows: (i) correlations among the three assets are low or even negative at the beginning of our sample but subsequently increase, and the change in the patterns becomes most pronounced after decisive structural breaks take place (breaks occur during the 2006–2009 period at different dates for specific asset pairs); (ii) correlations before the 2007–2008 crisis exhibit different patterns at different investment horizons; (iii) during and after the crisis, the correlations exhibit large swings and their differences at shorter and longer investment horizons become negligible. This finding indicates vanishing potential for risk diversification based on these assets: after the structural change, gold, oil, and stocks could not be combined to yield effective risk diversification during the post-break period studied.

The chapter is organized as follows. In Sect.1.2, we introduce the theoretical framework for the wavelet methodology we use to perform our analysis. Our large data set is described in detail in Sect.1.3 with a number of relevant commentaries. We present our empirical results in Sect.1.4. Sect.1.5 briefly concludes.

1.2 Theoretical framework for the methodologies employed

In the following section, we introduce the methodologies employed. While standard approaches (e.g., DCC GARCH and realized volatility) allow us to study the covariance matrix solely in the time domain, we are interested in studying its time-frequency dynamics. In other words, we are interested in determining how the correlations vary over time and various investment horizons. We are able to do so by using the innovative time-frequency approach of wavelet analysis. Wavelets are a relatively new method in economics, despite their potential benefits to economists (Ramsey, 2002; Gençay et al, 2002).

We continue with a brief introduction of the methodologies used to estimate the dynamic correlations, namely: (i) the parametric DCC GARCH approach; (ii) nonparametric realized measures; and (iii) a time-frequency approach in the form of a wavelet analysis.

1.2.1 Time-varying correlations: DCC GARCH methodology

In this section, we introduce the Dynamic Conditional Correlation Generalized Autoregressive Conditional Heteroscedasticity (DCC GARCH) model for estimating dynamic correlations in a multivariate setting. The DCC GARCH was proposed by Engle (2002) and is a logical extension of Bollerslevs constant conditional correlation (CCC) model (Bollerslev, 1990), in which the volatilities of each asset were allowed to vary over time but the correlations were time invariant. The DCC version, however, also allows for dynamics in the correlations.

Engle (2002) defines the covariance matrix as:

$$H_t = D_t R_t D_t. \tag{1.1}$$

where R_t is the conditional correlation matrix and $D_t = diag\{\sqrt{h_{i,t}}\}$ is a diagonal matrix of time varying standard variation from the i-th univariate (G)ARCH(p,q) processes $h_{i,t}$. Parameter n represents the number of assets at time $t = 1,\ldots,T$. The correlation matrix is then given by the transformation

$$R_t = diag(\sqrt{q_{11,t}},\ldots,\sqrt{q_{nn,t}})Q_t diag(\sqrt{q_{11,t}},\ldots,\sqrt{q_{nn,t}}), \tag{1.2}$$

where $Q_t = (q_{ij,t})$ is

$$Q_t = (1-\alpha-\beta)\overline{Q} + \alpha\eta_{t-1}\eta'_{t-1} + \beta Q_{t-1}, \tag{1.3}$$

where $\eta_t = \varepsilon_{i,t}/\sqrt{h_{i,t}}$ denotes the standardized residuals from the (G)ARCH model, $\overline{Q} = T^{-1}\sum \eta_t \eta'_t$ is a $n \times n$ unconditional covariance matrix of η_t, and α and β are non-negative scalars such that $\alpha + \beta < 1$.

We estimate the DCC GARCH using the standard quasi-maximum likelihood method proposed by Engle (2002). Further, we assume Gaussian innovations. The DCC model can be estimated consistently by estimating the univariate GARCH models in the first stage and the conditional correlation matrix in the second stage. The parameters are also estimated in stages. This two-step approach avoids the dimensionality problem encountered in most multivariate GARCH models (Engle, 2002; Engle and Sheppard, 2001).[1] Furthermore, the DCC model is parsimonious and ensures that time-varying correlation matrices between the stock exchange returns are positive definite.

1.2.2 Time-varying correlations: Realized Volatility approach

Due to increased availability of high–frequency data, a simple technique for estimating the covariance matrix was recently developed. In contrast to the DCC GARCH,

[1] Bauwens and Laurent (2005) demonstrate that the one-step and two-step methods provide very similar estimates.

this method is non-parametric. It is based on the estimating the covariance matrix analogously to the realized variation by taking the outer product of the observed high-frequency return over the period considered. Following Andersen et al (2003) and Barndorff-Nielsen and Shephard (2004) we define the realized covariance over the time interval $[t-h,t]$ for $0 \leq h \leq t \leq T$ as

$$\widehat{RC}_{t,h} = \sum_{i=1}^{M} \mathbf{r}_{t-h+\left(\frac{i}{M}\right)h} \mathbf{r}'_{t-h+\left(\frac{i}{M}\right)h}, \tag{1.4}$$

where M denotes the number of observations in the interval $[t-h,t]$. Andersen et al (2003) and Barndorff-Nielsen and Shephard (2004) demonstrate that the *ex-post* realized covariance $\widehat{RC}_{t,h}$ is an unbiased estimator of the *ex-ante* expected covariation. Furthermore, given increasing sampling frequency, i.e. $h > 0$ and $M \to \infty$, the realized covariance is a consistent estimator of the covariation. In practice, we only observe discrete prices, hence discretization bias is unavoidable. More serious damage to the estimator is also caused by market microstructure effects such as the bid-ask bounce, price discreteness, and the bid-ask spread. The literature advises employing rather sparse sampling when applying the estimator in practice; however this entails discarding a large amount of the available data. Following the suggestion by Andersen and Benzoni (2007) to obtain the best trade-off between reduced bias and information loss, we use 5-minute data to calculate the realized covariances.[2] An Important assumption regarding the price processes is that the data are synchronized, which implies collecting the prices simultaneously. However, this is not an issue in our analysis, as all three examined assets are paired using equal time-stamps matching.

1.2.3 Time-frequency dynamics in correlations: Wavelet approach

As we are interested in how the correlations vary over time and at different investment horizons, we need to conduct a wavelet analysis that allows us to work simultaneously in the time and frequency domains. The DCC GARCH and realized volatility methods outlined above do not allow the researcher to extend the analysis to the frequency domain; hence we are only able to study the covariance matrix in the time domain.

Wavelet time–frequency domain analysis is very powerful tool when we expect changes in economic relationships such as structural breaks. Wavelet analysis can react to these changes because the wavelet transform uses a localized function with finite support for the decomposition – a wavelet. In contrast, when using a pure frequency approach, represented by the Fourier transform, one obtains information on all of the frequency components, but because the amplitude is fixed throughout

[2] This is the optimal sampling frequency determined based on the substantial research on the noise-to-signal ratio. The literature is well surveyed by Hansen and Lunde (2006), Bandi and Russell (2006), McAleer and Medeiros (2008), and Andersen and Benzoni (2007).

the period considered, the time information is completely lost. Thus, in the event of sudden changes in economic relationships or the presence of breaks during the period studied, one is unable to locate precisely where this change occurs. Additionally, due to the non-stationarity induced by such breaks, Fourier transform-based estimates may not be precise. Therefore, the wavelet transform has substantial advantages over the Fourier transform when the time series is non-stationary or is only locally stationary (Roueff and Sachs, 2011).

An important feature of wavelet analysis is the decomposition of the economic relationship into time-frequency components. Wavelet analysis often uses scale instead of frequency, as scale typically characterizes frequency bands. The set of wavelet scales can be further interpreted as investment horizons at which we can study the economic relationships separately. Thus, every scale describes the development of the economic relationship at a particular frequency while retaining the time dynamics. Subsequently, the wavelet decomposition generally provides a more complex picture compared to the time domain approach, which aggregates all investment horizons. Therefore, if we expect that economic relationships follow different patterns at various investment horizons, then a wavelet analysis can uncover interesting characteristics of the data that would otherwise remain hidden. An introduction to the wavelet methodology with a remarkable application to economics and finance is provided in Gençay et al (2002) and Ramsey (2002).

1.2.4 Wavelet transform

While we use a discrete version of the wavelet transform, we begin our introduction with the continuous wavelet transform (CWT), as it is the cornerstone of the wavelet methodology. Next, we continue by describing a special form of discrete wavelet transform named the "maximal overlap discrete wavelet transform" (MODWT). Following standard notation, we define the continuous wavelet transform $W_x(j,s)$ as a projection of a wavelet function[3] $\psi_{j,s}(\mathrm{t}) = \frac{1}{\sqrt{j}}\psi\left(\frac{t-s}{j}\right) \in L^2(\mathbb{R})$ onto the time series $x(t) \in L^2(\mathbb{R})$,

$$W_x(j,s) = \int_{\infty}^{\infty} x(t) \frac{1}{\sqrt{j}} \overline{\psi\left(\frac{t-s}{j}\right)} dt, \tag{1.5}$$

where s determines the position of the wavelet in time. The scaling, or dilatation parameter j controls how the wavelet is stretched or dilated. If the scaling parameter j is low (high), then the wavelet is more (less) compressed and able to detect high (low) frequencies. One of the most important conditions a wavelet must fulfill is the admissibility condition: $C_\psi = \int_0^\infty \frac{|\Psi(f)|^2}{f} df < \infty$, where $\Psi(f)$ is the Fourier transform of a wavelet $\psi(.)$. The decomposed time series $x(t)$ can be subsequently recovered using the wavelet coefficients as follows

[3] We use the least asymmetric wavelet with length L=8, denoted as LA(8).

$$x(t) = \frac{1}{C_\psi} \int_0^\infty \left[\int_{-\infty}^{\infty} W_x(j,s) \psi_{j,s}(t) ds \right] \frac{dj}{j^2}, \quad s>0. \tag{1.6}$$

Further, the continuous wavelet transform preserves the energy or variance of the analyzed time series; hence

$$x^2 = \frac{1}{C_\psi} \int_0^\infty \left[\int_{-\infty}^{\infty} |W_x(j,s)|^2 ds \right] \frac{dj}{j^2}. \tag{1.7}$$

Equation 1.7 is an important property that allows us to work with the wavelet variance, covariance and the wavelet correlation. For a more detailed introduction to continuous wavelet transform and wavelets, see Daubechies (1992), Chui (1992), and Percival and Walden (2000).

As we study discrete time series, we only require a limited number of scales, and some form of discretization is needed. The counterpart of the continuous wavelet transform in discrete time is the discrete wavelet transform,[4] which is a parsimonious form of the continuous transform, but it has some limiting properties that make its application to real time series relatively difficult. These limitations primarily concern the restriction of the sample size to the power of two and the sensitivity to the starting point of the transform. Therefore, in our analysis, we use a modified version of the discrete wavelet transform—MODWT—which has some advantageous properties that are summarized below.

In contrast to the DWT, the MODWT does not use downsampling, as a consequence the vectors of the wavelet coefficients at all scales have equal length, corresponding to the length of transformed time series. Thus, the MODWT is not restricted to sample sizes that are powers of two. However, the MODWT wavelet coefficients are no longer orthogonal to each other at any scale. Additionally, the MODWT is a translation-invariant type of transform; therefore, it is not sensitive to the choice of the starting point of the examined process. Both the DWT and MODWT wavelet and scaling coefficients can be used for energy decomposition and analysis of variance of a time series in the time-frequency domain, however Percival (1995) demonstrates the dominance of the MODWT estimator of variance over the DWT estimator. Furthermore, Serroukh et al (2000) analyze the statistical properties of the MODWT variance estimator for non-stationary and non-Gaussian processes and show its statistical properties. For additional details on the MODWT, see Mallat (1998) and Percival and Walden (2000).

1.2.5 Maximal overlap discrete wavelet transform

This section demonstrates an application of the pyramid algorithm to obtain the MODWT wavelet and scaling coefficients. The method is based on filtering time

[4] For a definition and detailed discussion of the discrete wavelet transform, see Mallat (1998), Percival and Walden (2000), and Gençay et al (2002).

series with MODWT wavelet filters; the output after filtering is then filtered again in a subsequent stage to obtain other wavelet scales.

Let us begin with the first stage. The wavelet coefficients at the first scale $(j = 1)$ are obtained via circular filtering of time series x_t using the MODWT wavelet and scaling filters $h_{1,l}$ and $g_{1,l}$ (Percival and Walden, 2000) :

$$W_x(1,s) \equiv \sum_{l=0}^{L-1} h_{1,l}x(s-l \; modN), \quad V_x(1,s) \equiv \sum_{l=0}^{L-1} g_{1,l}x(s-l \; modN). \tag{1.8}$$

The second step of the algorithm uses the scaling coefficients $V_x(1,s)$ instead of x_t. The wavelet and scaling filters have a width $L_j = 2^{j-1}(L-1)+1$; therefore, for the second scale, the length of the filter is $L_2 = 15$. After filtering, we obtain the wavelet coefficients at scale $j = 2$:

$$W_x(2,s) \equiv \sum_{l=0}^{L-1} h_{2,l}V_x(1,s-l \; modN), \quad V_x(2,s) \equiv \sum_{l=0}^{L-1} g_{2,l}V_x(1,s-l \; modN). \tag{1.9}$$

After the two steps of the algorithm we have two vectors of the MODWT wavelet coefficients at scale $j = 1$ and $j = 2$; $W_x(1,s), W_x(2,s)$ and one vector of the MODWT wavelet scaling coefficients at scale two $V_x(2,s)$, where $s = 0,1,\ldots,N-1$ is the same for all vectors. The vector $W_x(1,s)$ represents wavelet coefficients that reflect variations at the frequency band $f[1/4,1/2]$, $W_x(2,s)$: $f[1/8,1/4]$ and $V_x(2,s)$: $f[0,1/8]$.

The transfer function of the filter $h_l : l = 0,1,\ldots,L-1$, where L is the width of the filter, is denoted as $H(.)$. The pyramid algorithm exploits the fact that if we increase the width of the filter to $2^{j-1}(L-1)+1$, then the filter with the impulse response sequence in the form:[5]

$$\{h_0, \underbrace{0,\ldots,0}_{2^{j-1}-1 \text{ zeros}}, h_1, \underbrace{0,\ldots,0}_{2^{j-1}-1 \text{ zeros}}, h_{L-2}, \underbrace{0,\ldots,0}_{2^{j-1}-1 \text{ zeros}}, h_L\}, \tag{1.10}$$

has a transfer function defined as $H\left(2^{j-1}f\right)$. Using this feature of the filters, we can write the pyramid algorithm simply in the following form

$$W_x(j,s) \equiv \sum_{l=0}^{L-1} h_l V_x\left(j-1, s-2^{j-1}l \; modN\right), \; s = 0,1,\ldots,\mathrm{N}-1, \tag{1.11}$$

$$V_x(j,s) \equiv \sum_{l=0}^{L-1} g_l V_x\left(j-1, s-2^{j-1}l \; modN\right), \; s = 0,1,\ldots,\mathrm{N}-1, \tag{1.12}$$

---

[5] The number of zeros between filter coefficients is $2^{j-1}-1$, i.e., for the filter at the first stage, we have no zeros, and for the second stage there is just one zero between each coefficient; hence the width of the filter is 15.

where for the first stage we set $x = V_x(0,s)$. Thus, after performing the MODWT, we obtain $J^m \leq log_2(N)$ vectors of wavelet coefficients and one vector of scaling coefficients. The j-th level wavelet coefficients in vector $W_x(j,s)$ represent frequency bands $f[1/2^{j+1}, 1/2^j]$, while the j-th level scaling coefficients in vector $V_x(j,s)$ represent $f[0, 1/2^{j+1}]$. In the subsequent analysis of the wavelet correlations we apply the MODWT with the wavelet filter LA(8), with reflecting boundary conditions.

1.2.6 Wavelet correlation

Applying the wavelet transform allows for a scale-by-scale decomposition of a time series, where every scale represents an investment horizon. In the bivariate case, where we decompose two time series, we can study correlations at investment horizons represented by scales. The method is called wavelet correlation and offers an alternative means of studying the dependence between two time series, as it can uncover different dependencies across available scales.

When the MODWT is used, all vectors of the wavelet coefficients have the same length. Thus, for a time series x_t, $t = 1, 2, \ldots, N$, we obtain $j = 1, \ldots, J^m$ vectors of wavelet coefficients and one vector of scaling coefficients of length N. The maximal level of wavelet decomposition is denoted J^m, $J^m \leq log_2(N)$. The wavelet correlation $\rho_{xy}(j)$ between time series x_t and y_t at scale j is then defined as (Whitcher et al, 2000):

$$\rho_{xy}(j) = \frac{cov\left(W_x(j,s)W_y(j,s)\right)}{\left[Var\left(W_x(j,s)\right)var\left(W_y(j,s)\right)\right]^{\frac{1}{2}}} \equiv \frac{\gamma_{xy}(j)}{\nu_x(j)\,\nu_y(j)}, \qquad (1.13)$$

where $\nu_x^2(j)$ and $\gamma_{xy}(j)$ denote wavelet variance and covariance, respectively. Additional details on wavelet variance and covariance are provided in Appendix 1.6 and 1.7. The wavelet correlation estimator directly uses the definition of the wavelet correlation Eq.1.13, thus we can write:

$$\widehat{\rho}_{xy}(j) \equiv \frac{\widehat{\gamma}_{xy}(j)}{\widehat{\nu}_x(j)\widehat{\nu}_y(j)}, \qquad (1.14)$$

where $\widehat{\gamma}_{xy}(j)$ denotes the wavelet covariance estimator and $\widehat{\nu}_x(j)^2$ and $\widehat{\nu}_y(j)^2$ are estimators of wavelet variance at scale j for time series x_t and y_t. Whitcher et al (1999) established the central limit theorem for estimator Eq.1.14, as well as the approximate confidence intervals; empirical values are reported in Sect.1.4.1. Additional details on this topic can be found in Serroukh et al (2000).

1.3 Data

In the empirical section, we analyze the prices of gold, oil, and a representative U.S. stock market index, the S&P 500. The data set contains the tick prices of the S&P 500 and the futures prices of gold and oil, where we use the most active rolling contracts from the pit (floor traded) session. All of the assets are traded on the platforms of the Chicago Mercantile Exchange (CME).[6]

We restrict our study to the intraday 5-minute and daily data sampled during the business hours of the New York Stock Exchange (NYSE), as most of the liquidity of the S&P 500 comes from the period when the U.S. markets are open. The sample period runs from January 2, 1987 until December 31, 2012.[7] To synchronize the data, we employ Greenwich Mean Time (GMT) stamp matching. Further, we exclude transactions executed on Saturdays and Sundays, U.S. federal holidays, December 24 to 26, and December 31 to January 2, as the low activity on these days could lead to estimation bias. Therefore, we use data from 6472 trading days. Descriptive statistics of the intra-day and daily returns of the data that form our sample are presented in Tab.1.1. Overall, the statistics are standard with the remarkable exception of a very high excess kurtosis of 104.561 for oil. This is mainly a consequence of a single positive price change of 16.3% (January 19, 1991), when the worst deliberate environmental damage in history was caused by Iraqi leader Saddam Hussein, who ordered a large amount of oil to be spilled into the Persian Gulf (Khordagui and Al-Ajmi, 1993).

Figure 1.1 depicts the development of the prices of the three assets, in which several recessions and crisis periods can be detected. Following the National Bureau of Economic Research (NBER),[8] there were three recessions in the U.S. during the period studied: July 1990 to March 1991, March 2001 to November 2001, and December 2007 to June 2009. These recessions are highlighted by gray bands. Furthermore, black lines depict one-day crashes associated with large price drops. Specifically, Black Monday (October 19, 1987), the Asian crisis (October 27, 1997), the Russian ruble devaluation (August 17, 1998), the dot-com bubble burst (March 10, 2000), the World Trade Center attacks (September 11, 2001), the Lehman Brothers Holdings bankruptcy (September 15, 2008), and the Flash Crash (March 6, 2010). The largest one-day drops in the studied sample occurred on the following dates, with percentage declines given in parentheses: October 19, 1987 (20.47%), October 26, 1987 (8.28%), September 29, 2008 (8.79%), October 9, 2008 (7.62%), October 15, 2008 (9.03%), and December 1, 2008 (8.93%).

The above crashes differ in nature, and we discuss them briefly below. On Monday, October 19, 1987, known as Black Monday, stock markets around the world

---

[6] Oil (Light Crude) is traded on the New York Mercantile Exchange (NYMEX) platform, gold is traded on the Commodity Exchange, Inc. (COMEX), a division of NYMEX, and the S&P 500 is traded at the CME in Chicago. All data were acquired from Tick Data, Inc.

[7] The CME introduced the Globex(R) electronic trading platform in December 2006 and began to offer nearly continuous trading.

[8] US Business Cycle Expansions and Contractions, NBER, accessed April 5, 2013 (http://www.nber.org/cycles.html).

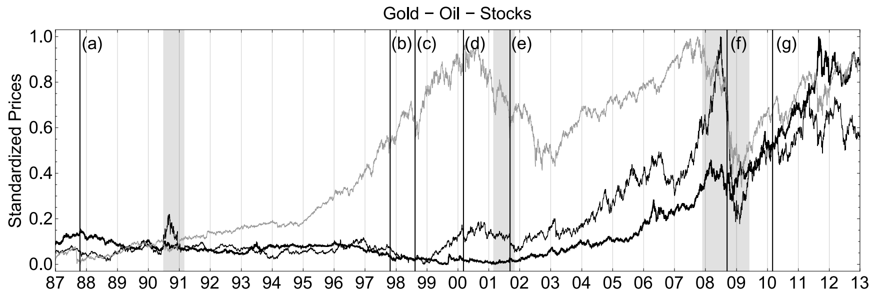

Fig. 1.1: Normalized prices of gold (thin black), oil (black), and stocks (gray). The figure highlights several important recession periods in gray (described in greater detail in the text), and crashes using black lines: (a) Black Monday; (b) the Asian crisis; (c) the Russian ruble devaluation; (d) the dot-com bubble burst; (e) the WTC 9/11 attacks; (f) the Lehman Brothers Holdings bankruptcy; and (g) the Flash Crash

dropped in a very short time and recorded the largest one-day crash in history. After this extreme event, many expected the most troubled years since the 1930s. Nevertheless, stock markets quickly recovered from the losses and closed 1987 in positive territory. There is still no consensus on the cause of the crash; potential reasons include illiquidity, program trading, overvaluation and market psychology.[9]

For many consecutive years stock markets did not record large shocks until 1996 when the Asian financial crisis emerged. Investors were leaving emerging overheated Asian shares that on October 27, 1997 resulted in a mini-crash of the U.S. markets. On August 17, 1998 the Russian government devalued the ruble, defaulted on domestic debt and declared a moratorium on payment to foreign creditors, which also caused an international crash. The 1996 and 1997 crashes are believed to be exogenous shocks to U.S. stock markets. The inflation of the so-called dot-com bubble emerged in the period 1997–2000, when several internet-based companies entered the markets and fascinated many investors confident in their future profits, while overlooking the companies' fundamental value. Ultimately, this resulted in a gradual collapse, or bubble burst, during the years 2000–2001. The World Trade Centre was attacked on September 11, 2001. Although markets recorded a sudden drop, the shock was exogenous and should not be attributed to internal market forces. The recent financial crisis of 2007–2008, also called as the global financial crisis (for a detailed treatment, see Bartram and Bodnar (2009)), was initiated by the bursting of the U.S. housing-market bubble. Consequently, in September and October 2008, stock markets experienced large declines. On May 6, 2010, financial markets witnessed the largest intraday drop in history known as the Flash Crash or The Crash of 2:45. The Dow Jones Industrial Average declined by approximately 1000 points (9%), but the loss was recovered within a few minutes. The crash was likely caused by high-frequency trading or large directional bets.

[9] For additional information on the crash, see Waldrop (1987) and Carlson (2007)

Table 1.1: Descriptive Statistics for high-frequency and daily gold, oil and, stock (S&P 500) returns over the sample period extending from January 2, 1987 until December 31, 2012

	High-frequency data			Daily data		
	Gold	Oil	Stocks	Gold	Oil	Stocks
Mean	1.00e-06	3.19e-06	-2.46e-06	2.22e-04	2.42e-04	2.70e-04
St. dev.	0.001	0.002	0.001	0.010	0.023	0.012
Skewness	-0.714	1.065	0.326	-0.147	-1.063	-0.392
Kurtosis	47.627	104.561	32.515	10.689	19.050	11.474
Minimum	-0.042	-0.045	-0.024	-0.077	-0.384	-0.098
Maximum	0.023	0.163	0.037	0.103	0.136	0.107

1.4 Empirical Analysis of the relationships among Gold, Oil, and Stocks

1.4.1 Dynamic correlations

Dynamic correlations for each pair of assets are depicted in Figs.1.2–1.4. Each figure consists of two panels that plot correlations obtained by the three methods described in Sect.1.2. The upper panels of the figures display realized volatility-based correlations computed on 5-minute returns for each day and daily correlations from the parametric DCC GARCH(1,1) estimates. The lower panels depict the evolution of time–frequency correlations obtained through a wavelet decomposition of 5-minute data.[10] The panel displays only four investment horizons as examples: 10 minutes, 40 minutes, 160 minutes and 1.6 days are depicted in the figures.

The correlations between asset pairs exhibit stable and similar patterns, where majority of time the correlations are low or even negative, until 2001 between gold and stocks, until 2004 between oil and stocks, and until 2005 between gold and oil. After these stable years, the pattern of the correlations fundamentally changes. The general pattern of dynamic correlations between the pairs of variables is the same regardless of what method is used. Nevertheless, there are noticeable differences. Correlations based on realized volatility provide very rough evidence. More contoured correlation patterns are inferred from the DCC GARCH method. The wavelet correlations illustrate the methods advantages over the two previous methods, as it allows us to observe individual correlation patterns for a number of investment horizons, providing time-frequency research output.[11]

In addition to the graphical illustration, the dynamic correlation results are summarized in Tabs. 1.2-1.4. The correlations for each asset pair are presented in individual tables containing the summarized correlations over a period of one year. The

[10] For the sake of clarity in the plot, we report monthly correlations, computed on monthly price time series.

[11] While the wavelet method is superior to the other two methods in terms of dynamic correlation analysis, we employ the other two methods as a benchmark.

tables have two main parts: the results in the left panels are based on high-frequency intraday data for different investment horizons ranging from 10 $(j=1)$ to 80 $(j=4)$ minutes, whereas the right panels contain daily correlations with investment horizons ranging from 2 to 32 days. Both panels also include a low-frequency component (approximately one year). With the aim of supporting the results, we compute confidence intervals around the reported point correlation estimates. The 95% confidence intervals of the estimates are nearly symmetric, with maximum values ranging from ± 0.014 for the first scale $j=1$ up to ± 0.04 for the last scale $j=4$.[12] Thus, based on the 95% confidence intervals, all reported correlation point estimates are statistically significant.

1.4.1.1 Gold–oil

The analysis of the intraday data for the gold-oil pair reveals a short period (1990–1991) of higher correlations, corresponding to the spike visible in Figs.1.2–1.4, which should be associated with the economic downturn in the U.S. from July 1990 to March 1991. During the period 1992–2005, the intraday correlations are remarkably low at short and longer horizons; see Tab.1.2. In 2006, a significant increase in correlation begins, reaching its maximum in 2012 at all investment horizons. In contrast to the period 1990–1991, the recent financial crisis changed the correlation structure of the gold and oil pair, indicating the existence of an important structural break in the correlation structure. This result is in line with the detected structural break on September 8, 2006 (Sect.1.4.2). Therefore, in terms of risk diversification, the situation changed dramatically for traders active at short-term investment horizons, as there is a significant increase in correlation after 2008 at all available investment horizons.

Dynamic correlations based on daily data reveal a more complex pattern. From 1987 until just before the global financial crisis erupted, correlations at diverse investment horizons seem quite heterogeneous (Tab.1.2). We observe very low correlations at short investment horizons measured in days, whereas at longer investment horizons of approximately one month, the correlations are higher. Beginning in 2008, the pattern changes significantly. Markets become quite homogeneous in perceptions of time because correlations at shorter and longer investment horizons become less diversified. Thus, the differences between short and long investment horizons diminish. One of the possible explanations is increased uncertainty in financial markets and poor economic performance in many developed countries during that period.

[12] For the sake of brevity, we do not report confidence intervals for all estimates. These results are available from the authors upon request.

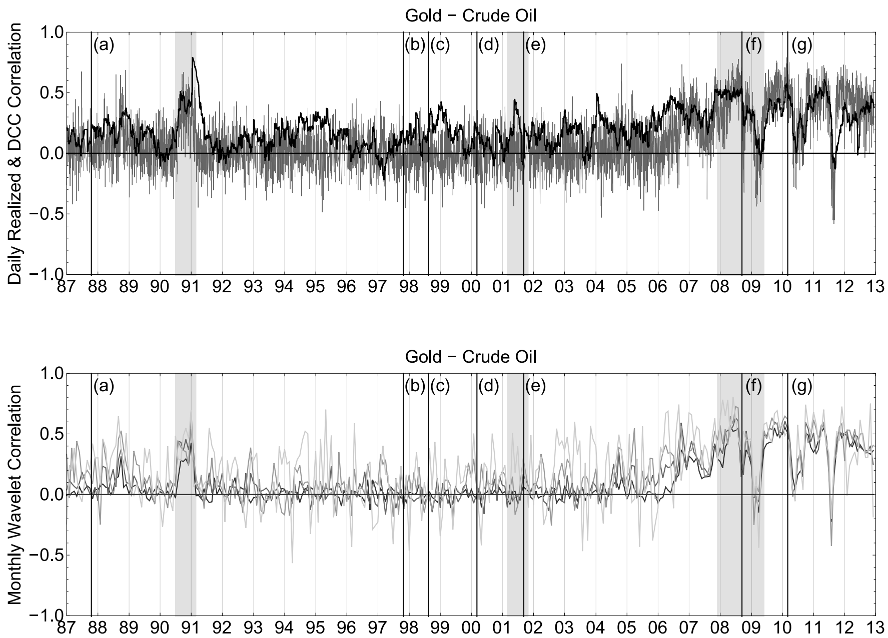

Fig. 1.2: Dynamics in gold-oil correlations. The upper plot of the panel contains the realized correlation for each day of the sample and daily correlations estimated from the DCC GARCH model. The lower plot contains time-frequency correlations based on the wavelet correlation estimates from high-frequency data for each month separately. We report correlation dynamics at 10-minute, 40-minute, 2.66 hour (approximate), and 1.6-day (approximate) horizons depicted by the thick black to thin black lines. The plots highlight several important recession periods in gray (described in greater detail in the text), and crashes using black lines: (a) Black Monday; (b) the Asian crisis; (c) the Russian ruble devaluation; (d) the dot-com bubble burst; (e) the WTC 9/11 attacks; (f) the Lehman Brothers Holdings bankruptcy; and (g) the Flash Crash

1.4.1.2 Gold–stocks and Oil–stocks

In comparison to the gold-oil pair, the gold-stocks and oil-stocks pairs provide a rather different picture (Tab.1.3 and Tab.1.4). During the period 1991–1992, negative correlations dominate, especially at longer horizons. The negative correlations are quite frequent for the two pairs, but they occur more frequently for the gold-stocks pair. Since 2001, the gold-stocks pair exhibits very rich correlation dynamics. The period of negative correlation begins in 2001, reaching its minimum in 2003, followed by a steady increase. After 2005, this pair exhibits significantly higher correlation, except for two short periods in 2008 and 2009.[13] In the 2012, we observe a significant increase in the correlation between gold and stocks at all available scales. There is an increase in magnitude that is three times larger relative to the previous

[13] On an annual basis, there was only a small decrease in 2011, as shown in Tab. 3.

year. This finding indicates a very limited possibility to diversify risk between stocks and gold in 2012.

The correlations of the oil-stocks pair also increased after the recent financial crisis began. Nevertheless, unlike the other two pairs, the correlation between oil and stocks before the crisis was considerably lower than after the crisis. This implies that the developments in 2008 had the strongest impact on the correlation structure of this pair. Further, from 2008 on, this pair has the highest correlation of the three examined pairs and highly homogeneous correlations at all scales. Therefore, after 2008 until the end of our sample, we only observe a negligible possibility for risk diversification in the sense of various investment horizons.

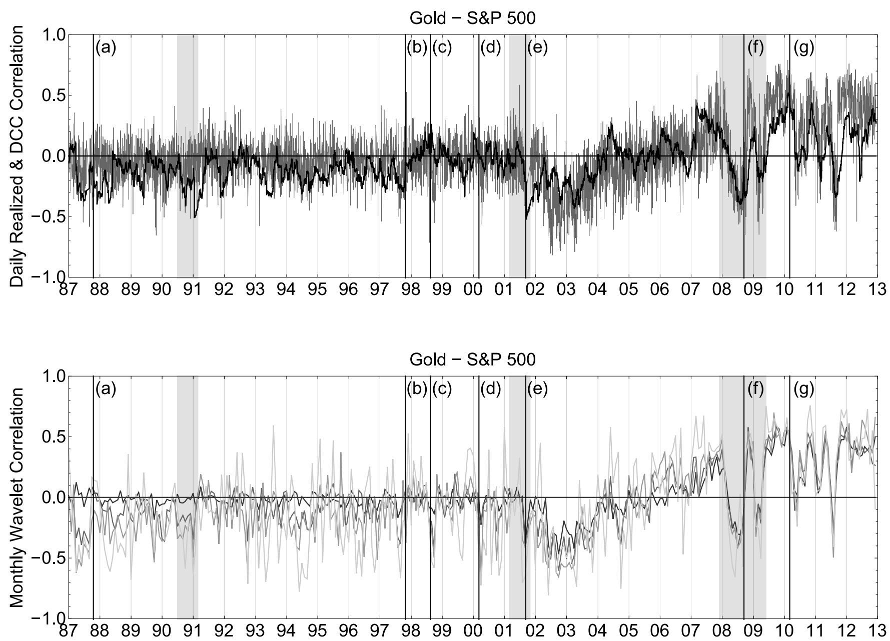

Fig. 1.3: Dynamics in gold-stocks correlations. The upper plot of the panel contains the realized correlation for each day of the sample and daily correlations estimated from the DCC GARCH model. The lower plot contains time-frequency correlations based on the wavelet correlation estimates from high-frequency data for each month separately. We report correlation dynamics at 10-minute, 40-minute, 2.66 hour (approximate), and 1.6-day (approximate) horizons depicted by the thick black to thin black lines. The plots highlight several important recession periods in gray (described in greater detail in the text), and crashes using black lines: (a) Black Monday; (b) the Asian crisis; (c) the Russian ruble devaluation; (d) the dot-com bubble burst; (e) the WTC 9/11 attacks; (f) the Lehman Brothers Holdings bankruptcy; and (g) the Flash Crash

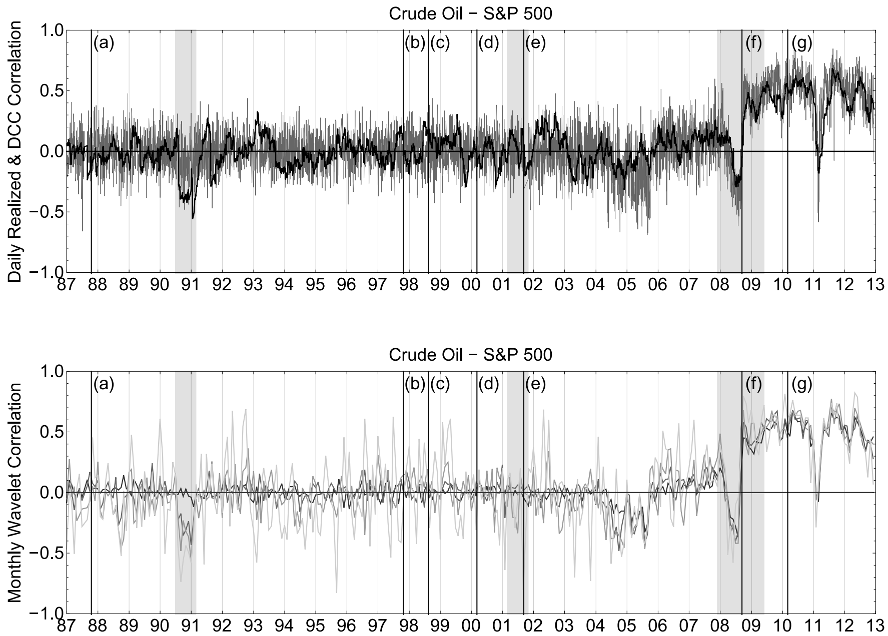

Fig. 1.4: Dynamics in oil-stocks correlations. The upper plot of the panel contains the realized correlation for each day of the sample and daily correlations estimated from the DCC GARCH model. The lower plot contains time-frequency correlations based on the wavelet correlation estimates from high-frequency data for each month separately. We report correlation dynamics at 10-minute, 40-minute, 2.66 hour (approximate), and 1.6-day (approximate) horizons depicted by the thick black to thin black lines. The plots highlight several important recession periods in gray (described in greater detail in the text), and crashes using black lines: (a) Black Monday; (b) the Asian crisis; (c) the Russian ruble devaluation; (d) the dot-com bubble burst; (e) the WTC 9/11 attacks; (f) the Lehman Brothers Holdings bankruptcy; and (g) the Flash Crash

1.4.2 Risk diversification

A wavelet methodology allows us to decompose mutual dependencies into different investment horizons; subsequently, we are able to generalize inferences related to risk diversification. When correlations are heterogeneous in their magnitudes at various investment horizons, market participants are able to diversify risk across these investment horizons, represented by scales. However, negligible or even no differences in correlation magnitudes at different investment horizons prevent effective risk diversification. For our set of assets, there was room for risk diversification until 2001; whereas problems with risk diversification arise after 2001.

Structural breaks cause important change with respect to the heterogeneity of correlations. In our analysis, we test for structural breaks in the correlations by employing the supF test (Hansen, 1992; Andrews, 1993; Andrews and Ploberger,

Table 1.2: Time-frequency correlation estimates for the gold–oil pair. The high-frequency set contains Wavelet correlation estimates based on high-frequency data. The daily set contains Wavelet correlation estimates based on daily data

	High-frequency data (minutes)					Daily data (days)					
	10	20	40	80	160-y.	2	4	8	16	32	64- y.
1987	0.02	0.03	0.08	0.15	0.80	0.12	0.00	0.14	0.07	-0.13	0.77
1988	0.11	0.19	0.19	0.23	0.42	0.23	0.25	0.32	0.38	0.55	0.93
1989	0.02	0.03	0.06	0.06	0.53	0.02	0.11	0.05	0.09	0.74	-0.14
1990	0.16	0.27	0.30	0.29	0.43	0.51	0.47	0.40	0.33	0.76	0.69
1991	0.21	0.32	0.31	0.32	0.63	0.01	0.00	0.47	0.48	-0.31	-0.47
1992	0.02	0.08	0.03	0.01	-0.56	0.06	0.01	-0.18	0.25	-0.05	-0.41
1993	0.00	0.01	0.04	0.02	0.60	0.12	0.04	0.20	0.30	-0.26	-0.77
1994	0.02	0.02	0.03	0.03	-0.13	0.16	0.37	0.22	-0.09	-0.28	0.33
1995	0.01	0.00	0.03	0.07	0.05	0.23	0.17	0.07	-0.02	0.16	0.39
1996	0.01	0.02	0.00	0.04	-0.62	-0.09	-0.03	0.13	-0.34	-0.31	-0.68
1997	0.00	-0.01	0.00	0.06	0.33	0.00	-0.22	0.04	-0.13	0.09	0.57
1998	0.00	-0.02	-0.01	-0.01	0.65	0.14	0.28	0.40	0.21	0.65	0.18
1999	0.01	0.01	-0.01	0.02	-0.58	-0.02	0.12	0.31	-0.17	0.17	-0.80
2000	0.00	0.00	0.01	0.07	-0.68	0.16	0.03	0.32	-0.12	0.01	0.44
2001	0.00	0.01	0.01	0.02	-0.83	0.23	0.04	0.11	-0.25	-0.10	0.10
2002	-0.01	-0.01	0.04	0.07	0.62	0.10	0.03	-0.17	0.08	-0.64	0.86
2003	0.01	0.02	0.04	0.06	0.68	0.24	0.05	-0.08	0.12	0.47	0.54
2004	0.04	0.08	0.10	0.11	0.40	0.17	0.38	0.23	0.13	-0.58	-0.74
2005	0.01	0.07	0.09	0.11	-0.42	0.08	0.07	0.22	0.42	0.27	0.40
2006	0.11	0.17	0.30	0.35	0.74	0.37	0.53	0.47	0.58	0.57	0.92
2007	0.26	0.30	0.33	0.35	0.29	0.49	0.38	0.07	0.41	0.42	0.48
2008	0.32	0.35	0.39	0.39	0.74	0.44	0.45	0.55	0.67	0.41	0.27
2009	0.19	0.21	0.22	0.22	-0.21	0.19	0.20	0.53	-0.03	-0.12	0.45
2010	0.33	0.34	0.36	0.37	-0.30	0.29	0.35	0.48	0.57	0.07	-0.37
2011	0.26	0.27	0.31	0.29	0.22	0.20	0.18	0.20	0.37	0.62	0.72
2012	0.40	0.42	0.42	0.41	-0.36	0.37	0.40	0.63	0.43	-0.19	0.71

1994) with p-values computed based on Hansen (1997); the results are summarized in Tab.1.5. An illustrative example of a pre-structural break period is the gold-oil pair, with the break detected on September 8, 2006. We can observe a significant increase in the overall correlation estimated by the DCC GARCH during the periods 1994–1996 and 1998–2000. The DCC GARCH estimates aggregate the correlation over all investment horizons. However, using wavelet correlations, we obtain additional information that this increase might be caused particularly by the long-term correlations, as the correlations at short investment horizons are close to zero. Similar patterns are observed for the gold-stocks and oil-stocks pairs, for which structural breaks were detected on May 5, 2009 and September 26, 2008, respectively.

Thus, we observe that the correlations between asset pairs were very heterogeneous across investment horizons before the structural break. Conversely, after the structural break, the correlation pattern became mostly homogeneous, which implies that gold, oil and stocks could no longer be simultaneously included in a single portfolio for risk diversification purposes. This finding contradicts the results of

Table 1.3: Time-frequency correlation estimates for the gold–stocks pair. The high-frequency set contains Wavelet correlation estimates based on high-frequency data. The daily set contains Wavelet correlation estimates based on daily data

	High-frequency data (minutes)					Daily data (days)					
	10	20	40	80	160-y.	2	4	8	16	32	64- y.
1987	0.05	-0.04	-0.11	-0.22	-0.54	-0.22	-0.22	-0.24	-0.39	-0.58	0.64
1988	-0.02	-0.05	-0.14	-0.22	-0.21	-0.25	0.08	-0.06	-0.17	-0.12	-0.54
1989	-0.03	-0.11	-0.19	-0.15	-0.59	0.03	-0.26	-0.23	0.06	-0.67	-0.92
1990	-0.04	-0.15	-0.20	-0.25	-0.77	-0.32	-0.33	-0.28	-0.10	-0.36	-0.84
1991	-0.01	-0.07	-0.10	-0.09	-0.55	-0.16	-0.14	0.11	0.18	0.14	-0.59
1992	-0.03	-0.04	-0.03	-0.10	0.52	0.01	-0.01	-0.21	-0.28	0.37	0.21
1993	-0.02	-0.06	-0.10	-0.13	0.49	-0.24	-0.15	-0.23	0.03	-0.12	0.43
1994	-0.02	-0.10	-0.17	-0.21	0.41	-0.26	-0.18	0.05	0.00	0.38	0.31
1995	-0.01	-0.04	0.01	-0.04	-0.37	-0.17	-0.06	-0.02	0.01	-0.39	0.71
1996	-0.04	-0.12	-0.08	-0.13	-0.26	-0.20	-0.27	-0.24	0.47	0.57	-0.77
1997	-0.03	-0.06	-0.07	-0.11	-0.49	-0.15	-0.17	-0.26	-0.05	0.03	-0.93
1998	-0.05	-0.07	-0.13	-0.11	0.80	-0.03	0.17	0.28	-0.05	0.49	0.43
1999	-0.01	-0.01	-0.04	0.01	0.54	-0.03	0.10	0.05	0.15	0.45	-0.75
2000	-0.03	-0.07	-0.10	-0.20	0.54	-0.03	0.00	0.24	0.32	-0.25	-0.80
2001	-0.01	-0.01	0.00	0.01	-0.49	-0.24	-0.11	0.07	0.04	0.16	0.43
2002	-0.27	-0.34	-0.38	-0.37	-0.58	-0.21	-0.24	-0.37	-0.28	-0.34	-0.66
2003	-0.26	-0.35	-0.38	-0.42	0.46	-0.42	-0.12	-0.07	-0.51	-0.49	0.18
2004	-0.07	-0.09	-0.08	-0.09	0.65	0.03	0.14	0.38	0.14	0.17	0.29
2005	-0.02	-0.02	0.02	0.00	0.08	-0.08	0.11	0.09	-0.02	0.40	0.13
2006	0.05	0.11	0.17	0.20	0.30	0.10	-0.01	0.20	0.34	0.65	0.16
2007	0.20	0.26	0.29	0.27	-0.18	0.39	0.28	0.42	0.42	0.85	0.39
2008	0.11	0.14	0.10	0.09	0.87	-0.03	-0.16	-0.09	-0.16	-0.68	-0.89
2009	0.14	0.13	0.15	0.17	-0.17	0.01	-0.05	0.28	0.31	-0.05	-0.38
2010	0.25	0.25	0.28	0.29	0.06	0.14	0.29	0.46	0.38	-0.08	-0.20
2011	0.13	0.14	0.18	0.13	-0.40	-0.17	-0.18	-0.08	-0.04	0.20	0.49
2012	0.40	0.39	0.37	0.38	0.67	0.42	0.26	0.62	0.58	0.06	0.05

Baur and Lucey (2010), who find gold to be a good hedge against stocks and therefore a safe haven during financial market turmoil. However, our result is in line with the argument of Bartram and Bodnar (2009) that diversification provided little help for investors during the financial crisis.

The change in the correlation structure described above can also be attributed to changes in investors' beliefs,[14] which become mostly homogeneous across investment horizons after the structural break. The homogeneity can be partially induced by broader uncertainty regarding financial markets pricing fundamentals.[15] Investors tendency to favor more aggressive strategies may be one of the reasons that we observe increased homogeneity in the correlations across investment hori-

[14] Additional information on the role of investors' beliefs can be found in Ben-David and Hirshleifer (2012).

[15] Connolly et al (2007) study the importance of time-varying uncertainty on asset correlation that subsequently influences the availability of diversification benefits.

Table 1.4: Time-frequency correlation estimates for the oil–stocks pair. The high-frequency set contains Wavelet correlation estimates based on high-frequency data. The daily set contains Wavelet correlation estimates based on daily data

	High-frequency data (minutes)					Daily data (days)					
	10	20	40	80	160-y.	2	4	8	16	32	64- y.
1987	0.03	0.01	0.05	0.04	-0.64	-0.11	0.21	-0.07	-0.09	-0.03	0.66
1988	0.03	-0.03	-0.05	-0.11	0.30	-0.06	0.13	-0.09	-0.15	-0.30	-0.74
1989	0.01	0.00	0.03	-0.02	0.13	-0.08	0.12	-0.06	-0.10	-0.74	0.06
1990	-0.02	-0.12	-0.18	-0.20	-0.54	-0.38	-0.46	-0.62	-0.25	0.04	-0.84
1991	-0.04	-0.10	-0.17	-0.19	-0.49	-0.06	-0.06	0.38	0.34	-0.51	0.58
1992	0.03	0.01	0.04	-0.02	-0.51	0.08	0.04	0.20	-0.20	-0.52	0.50
1993	0.00	0.00	-0.01	-0.02	0.73	-0.05	-0.11	0.24	0.42	0.10	-0.63
1994	0.00	0.01	-0.03	-0.04	-0.77	-0.23	-0.05	0.01	0.13	-0.47	-0.58
1995	-0.02	0.01	0.02	0.00	-0.47	-0.05	0.04	0.06	0.37	-0.14	-0.20
1996	0.00	0.00	0.00	-0.03	-0.02	-0.02	0.05	-0.18	-0.15	-0.38	0.56
1997	0.00	0.00	0.07	0.08	-0.61	-0.15	0.00	0.02	0.08	0.19	-0.50
1998	0.00	-0.02	0.02	0.02	0.64	0.04	0.13	0.15	0.21	0.52	-0.79
1999	-0.01	0.02	0.03	0.00	-0.94	-0.03	0.01	0.09	0.17	0.73	0.99
2000	0.02	-0.01	-0.02	-0.05	-0.18	-0.11	-0.09	0.07	-0.11	-0.53	-0.24
2001	0.01	0.04	0.03	-0.05	0.51	-0.12	-0.04	0.08	0.03	0.84	0.81
2002	-0.01	-0.01	-0.02	-0.03	-0.70	0.17	0.19	0.41	0.24	0.36	-0.53
2003	-0.01	-0.03	-0.04	-0.05	0.57	-0.24	0.08	-0.49	-0.36	-0.58	-0.54
2004	-0.07	-0.15	-0.18	-0.25	0.13	-0.13	0.01	0.05	-0.08	-0.68	-0.71
2005	-0.16	-0.19	-0.17	-0.18	0.00	-0.09	0.14	0.04	-0.46	-0.20	0.32
2006	0.04	0.07	0.09	0.10	-0.08	0.07	0.07	0.08	0.24	0.46	-0.12
2007	0.09	0.13	0.13	0.12	-0.63	0.17	0.07	-0.04	0.04	0.18	0.74
2008	0.26	0.27	0.31	0.33	0.70	0.42	0.32	0.09	0.05	0.12	-0.47
2009	0.42	0.46	0.48	0.50	0.92	0.53	0.28	0.61	0.10	-0.05	0.51
2010	0.57	0.59	0.58	0.62	0.69	0.70	0.71	0.51	0.58	0.91	0.86
2011	0.50	0.53	0.53	0.56	0.32	0.53	0.57	0.62	0.53	-0.03	0.74
2012	0.49	0.48	0.47	0.46	0.26	0.52	0.54	0.74	0.53	0.12	0.30

zons. Furthermore, the homogeneity in correlations may have been increased by the introduction of completely electronic trading on exchange platforms in 2005, which was accompanied by an increased volume of automatic trading.

Table 1.5: Values of the supF test with corresponding p-values. The break dates dividing the period into pre-break and post-break. The full period covers January 2, 1987 to December 31, 2012

	gold–oil	gold–stocks	oil –stocks
Date of the break	September 8, 2006	May 5, 2009	September 26, 2008
supF	3390.3	2544.3	7284.9
p-value	$< 2.2e\text{-}16$	$< 1.1e\text{-}16$	$< 2.2e\text{-}16$

1.5 Conclusions

In this chapter, we study dynamic co-movements among key traded assets by employing the realized volatility and DCC GARCH approaches as a benchmark and the wavelet methodology, a novel time-frequency approach. In terms of the dynamic method, the wavelet-based correlation analysis enables us to analyze co-movements among assets, not only from a time series perspective, but also from the investment horizon perspective. Thus, we are able to provide unique evidence on how correlations among major assets vary over time and at different investment horizons. We analyze dynamic correlations in the prices of gold, oil, and a broad U.S. stock market index, the S&P 500, over 26 years from January 2, 1987 until December 31, 2012. The analysis is performed on both intra-day and daily data.

Our findings suggest that the wavelet analysis outperforms the standard benchmark approaches. Further, it offers a crucial message based on the evidence of very different patterns in linkages among assets over time. During the period before the pairs of assets suffered from structural breaks, our results revealed very low, even negative, but heterogeneous correlations for all pairs. After the breaks, the correlations for all pairs increased on average, but their magnitudes exhibited large positive and negative swings. Surprisingly, despite this strongly varying behavior, the correlations between pairs of assets became homogeneous and did not differ across distinct investment horizons. A strong implication emerges. Prior to the structural break, it was possible to use all three assets in a well-diversified portfolio. However, after the structural changes occurred, gold, oil, and stocks could not be used in conjunction for risk diversification purposes during the post-break period studied.

Acknowledgements

We benefited from valuable comments we received from Abu Amin, Ladislav Krištoufek, Brian Lucey, Paresh Narayan, Lucjan Orlowski, Perry Sadorsky, Yi-Ming Wei, and Yue-Jun Zhang. The usual disclaimer applies. The support from the Czech Science Foundation (GAČR) under Grants GA13-24313S and GA14-24129S is gratefully acknowledged.

References

Aggarwal R, Lucey BM (2007) Psychological barriers in gold prices? Review of Financial Economics 16(2):217–230

Aguiar-Conraria L, Martins M, Soares MJ (2012) The yield curve and the macroeconomy across time and frequencies. Journal fo Economic Dynamics and Control 36:1950–1970

Andersen T, Benzoni L (2007) Realized volatility. In: Andersen T, Davis R, Kreiss J, Mikosch T (eds) Handbook of Financial Time Series, Springer Verlag

Andersen T, Bollerslev T, Diebold F, Labys P (2003) Modeling and forecasting realized volatility. Econometrica 71(2):579–625

Andrews DW (1993) Tests for parameter instability and structural change with unknown change point. Econometrica 61:821—856

Andrews DW, Ploberger W (1994) Optimal tests when a nuisance parameter is present only under the alternative. Econometrica 62:1383–1414

Bandi F, Russell J (2006) Volatility. In J. Birge and V. Linetsky (Eds.), Handbook of Financial Engineering. Elsevier

Barndorff-Nielsen O, Shephard N (2004) Econometric analysis of realized covariation: High frequency based covariance, regression, and correlation in financial economics. Econometrica 72(3):885–925

Bartram SM, Bodnar GM (2009) No place to hide: The global crisis in equity markets in 2008/2009. Journal of International Money and Finance 28(8):1246–1292

Baur DG, Lucey BM (2010) Is gold a hedge or a safe haven? an analysis of stocks, bonds and gold. Financial Review 45(2):217–229

Bauwens L, Laurent S (2005) A new class of multivariate skew densities, with application to generalized autoregressive conditional heteroscedasticity models. Journal of Business & Economic Statistics 23(3):346–354

Ben-David I, Hirshleifer D (2012) Are investors really reluctant to realize their losses? trading responses to past returns and the disposition effect. Review of Financial Studies 25(8):2485–2532

Bollerslev T (1990) Modelling the coherence in short-run nominal exchange rates: A multivariate generalized arch approach. Review of Economics and Statistics 72(3):498–505

Büyükşahin B, Robe MA (2013) Speculation, commodities and cross-market linkages. Journal of International Money and Finance

Carlson M (2007) A brief history of the 1987 stock market crash with a discussion of the federal reserve response. Divisions of Research & Statistics and Monetary Affairs, Federal Reserve Board

Cashin P, McDermott C, Scott A (1999) The myth of co-moving commodity prices. IMF Working Paper WP/99/169 International Monetary Fund, Washington

Cashin P, McDermott CJ, Scott A (2002) Booms and slumps in world commodity prices. Journal of Development Economics 69(1):277–296

Chui C (1992) An Inroduction to Wavelets. Academic Press, New York

Connolly RA, Stivers C, Sun L (2007) Commonality in the time-variation of stock–stock and stock–bond return comovements. Journal of Financial Markets 10(2):192–218

Daubechies I (1992) Ten lectures on wavelets. SIAM

Engle R (2002) Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. Journal of Business & Economic Statistics 20(3):339–350

Engle RF, Sheppard K (2001) Theoretical and empirical properties of dynamic conditional correlation multivariate garch. Tech. rep., National Bureau of Economic Research

Faÿ G, Moulines E, Roueff F, Taqqu M (2009) Estimators of long-memory: Fourier versus wavelets. Journal of Econometrics 151(2):159–177

Fratzscher M, Schneider D, Van Robays I (2013) Oil prices, exchange rates and asset prices. CESifo Working Paper No 4264

Gadanecz B, Jayaram K (2009) Measures of financial stability–a review. Bank for International Settlements, IFC Bulletin 3

Gallegati M, Gallegati M, Ramsey JB, Semmler W (2011) The us wage phillips curve across frequencies and over time. Oxford Bulletin of Economics and Statistics 74(4):489–508

Gençay R, Selçuk F, Whitcher B (2002) An Introduction to Wavelets and Other Filtering Methods in Finance and Economics. Academic Press

Graham M, Kiviaho J, Nikkinen J (2013) Short-term and long-term dependencies of the s&p 500 index and commodity prices. Quantitative Finance 13(4):583–592

Hamilton J (2009) Causes and consequences of the oil shock of 2007-08. Brookings Papers in Economic Activity 40(1):215–283

Hansen BE (1992) Tests for parameter instability in regressions with i (1) processes. Journal of Business & Economic Statistics 10:321–335

Hansen BE (1997) Approximate asymptotic p values for structuras-change tests. Journal of Business & Economic Statistics 15(1):60–67

Hansen P, Lunde A (2006) Realized variance and market microstructure noise. Journal of Business and Economic Statistics 24(2):127–161

Khordagui H, Al-Ajmi D (1993) Environmental impact of the gulf war: an integrated preliminary assessment. Environmental Management 17(4):557–562

Mallat S (1998) A wavelet tour of signal processing. Academic Press

Markowitz H (1952) Portfolio selection. The Journal of Finance 7(1):77–91

Marshall JF (1994) The role of the investment horizon in optimal portfolio sequencing (an intuitive demonstration in discrete time). Financial Review 29(4):557–576

McAleer M, Medeiros MC (2008) Realized volatility: A review. Econometric Reviews 27(1-3):10–45

Percival D, Walden A (2000) Wavelet Methods for Time Series Analysis. Cambridge University Press, Cambridge

Percival DB (1995) On estimationof the wavelet variance. Biometrika 82:619–631

Pindyck RS, Rotemberg JJ (1990) The excess co-movement of commodity prices. The Economic Journal 100(403):1173–89

Ramsey JB (2002) Wavelets in economics and finance: Past and future. Studies in Nonlinear Dynamics & Econometrics 6(3)

Roueff F, Sachs R (2011) Locally stationary long memory estimation. Stochastic Processes and their Applications 121(4):813 – 844

Samuelson PA (1989) The judgment of economic science on rational portfolio management: indexing, timing, and long-horizon effects. The Journal of Portfolio Management 16(1):4–12

Serroukh A, Walden AT, Percival DB (2000) Statistical properties and uses of the wavelet variance estimator for the scale analysis of time series. Journal of the American Statistical Association 95:184–196

Vacha L, Barunik J (2012) Co-movement of energy commodities revisited: Evidence from wavelet coherence analysis. Energy Economics 34(1):241–247

Waldrop MM (1987) Computers amplify black monday: The sudden stock market decline raised questions about the role of computers; they may not have actually caused the crash, but may well have amplified it. Science (New York, NY) 238(4827):602

Whitcher B, Guttorp P, Percival DB (1999) Mathematical background for wavelets estimators for cross covariance and cross correlation. Tech Rep 38, Natl Res Cent for stat and the Environ

Whitcher B, Guttorp P, Percival DB (2000) Wavelet analysis of covariance with application to atmosferic time series. Journal of Geophysical Research 105:941–962

Appendix

1.6 Wavelet variance

The variance of a time series can be decomposed into its frequency components, which are called scales in the wavelet methodology. Using wavelets, we can identify the portion of variance attributable to a specific frequency band of the examined time series. In this section, we demonstrate how to estimate wavelet variance and demonstrate that the summation of all of the components of wavelet variance yields the variance of the time series.

Let us suppose a real-valued stochastic process x_i, $i= 1,\dots,N$, whose $\frac{L}{2}$th backward difference is a covariance stationary stochastic process with mean zero. Then, the sequence of the MODWT wavelet coefficients $W_x(j,s)$, unaffected by the boundary conditions, for all scales $j = 1,2,\dots,J^m$ is also a stationary process with mean zero. As we use the least asymmetric wavelet of length $L = 8$, we can expect stationarity of the MODWT wavelet coefficients. Following Percival (1995), we define the wavelet variance at scale j as the variance of wavelet coefficients at scale j as:

$$\nu_x(j)^2 = var\left(W_x(j,s)\right). \tag{1.15}$$

For coefficients unaffected by the boundary conditions, which are defined for each scale separately $M_j = N - L_j + 1 > 0$, the unbiased estimator of wavelet variance at scale j reads:

$$\nu_x(j)^2 = \frac{1}{M_j}\sum_{s=L_j-1}^{N-1} W_x(j,s)^2. \tag{1.16}$$

As the variance of a covariance stationary process x_i is equal to the integral of the spectral density function $S_x(.)$, the wavelet variance at a scale j is the variance of the wavelet coefficients $W_x(j,s)$ with spectral density function $S_x(j)(.)$:

$$\nu_x(j)^2 = \int_{-1/2}^{1/2} S_x(j)(f)df = \int_{-1/2}^{1/2} \mathscr{H}_j(f) S_x(j)(f)df, \tag{1.17}$$

where $\mathscr{H}_j(f)$ is the squared gain function of the wavelet filter h_j (Percival and Walden, 2000). As the variance of a process x_i is the sum of the contributions of the wavelet variances at all scales we can write:

$$var(x) = \sum_{j=1}^{\infty} \nu_x(j)^2. \tag{1.18}$$

In case we have only a finite number of scales, we have to add also variance of the scaling coefficients vector; thus we can write:

$$var(x) = \int_{-1/2}^{1/2} S_x(f)df = \sum_{j=1}^{J^m} \nu_x(j)^2 + var\left(V_x(J^m, s)\right). \tag{1.19}$$

1.7 Wavelet covariance

The wavelet covariance of two processes x_t and y_t is estimated in a similar manner as the wavelet variance. As a first step, we perform the MODWT to obtain vectors of wavelet and scaling coefficients at all scales $j = 1, 2, \ldots, J^m$. While we use the LA(8) wavelet with length $L = 8$, we can use non-stationary processes, which are stationary after the d-th difference, where $d \leq L/2$. The wavelet covariance of x_t and y_t at scale j is defined as:

$$\gamma_{xy}(j) = Cov\left(W_x(j,s), W_y(j,s)\right). \tag{1.20}$$

Taking into consideration the MODWT wavelet coefficients unaffected by boundary conditions denoted $M_j = N - L_j + 1 > 0$, then for processes x_t and y_t defined above, the estimator of the wavelet covariance at a scale j is defined as

$$\widehat{\gamma}_{xy}(j) = \frac{1}{M_j} \sum_{s=L_j-1}^{N-1} W_x(j,s) W_y(j,s), \tag{1.21}$$

When processes x_t and y_t are Gaussian, the MODWT estimator of wavelet covariance is unbiased and asymptotically normally distributed (Whitcher et al, 1999). When we have an infinite time series, the number of available scales goes to infinity, $J^m \to \infty$, then the sum of all available wavelet covariances $\gamma_{xy}(j)$ yields the covariance of x_t and y_t:

$$Cov(x_t, y_t) = \sum_{j=1}^{\infty} \gamma_{xy}(j). \tag{1.22}$$

For a finite real time series, the number of scales is limited by $J^m \leq log_2(T)$, the covariance of x_t and y_t is a sum of covariances of the MODWT wavelet coefficients $\gamma_{xy}(j)$ at all scales $j = 1, 2, \ldots, J^m$ and the covariance of the scaling coefficients $V_x(J,s)$ at scale J^m:

$$Cov(x_t, y_t) = Cov(V_x(J^m, s), V_y(J^m, s)) + \sum_{j=1}^{J^m} \gamma_{xy}(j). \tag{1.23}$$

1.8 MODWT wavelet filters

Let us introduce the MODWT scaling and wavelet filters g_l and h_l, $l = 0, 1, \ldots, L-1$, where L denotes the length of the wavelet filter. For example, the Least Asymmetric (LA8) wavelet filter has length $L = 8$ (Daubechies, 1992). Generally, the scaling filter is a low-pass filter whereas the wavelet filter is a high-pass filter. There are three basic properties that both MODWT filters must satisfy. Let us describe these properties for the MODWT wavelet filter

$$\sum_{l=0}^{L-1} h_l = 0, \quad \sum_{l=0}^{L-1} h_l^2 = 1/2, \quad \sum_{l=-\infty}^{\infty} h_l h_{l+2N} = 0, \quad N \in \mathbb{Z}_N \tag{1.24}$$

and for the MODWT scaling filter

$$\sum_{l=0}^{L-1} g_l = 1, \quad \sum_{l=0}^{L-1} g_l^2 = 1/2, \quad \sum_{l=-\infty}^{\infty} g_l g_{l+2N} = 0, \quad N \in \mathbb{Z}_N. \tag{1.25}$$

The transfer function of a MODWT filter $\{h_l\}$ at frequency f is defined via the Fourier transform as

$$H(f) = \sum_{l=-\infty}^{\infty} h_l e^{-i2\pi fl} = \sum_{l=0}^{L-1} h_l e^{-i2\pi fl} \tag{1.26}$$

with the squared gain function defined as: $\mathscr{H}(f) = |H(f)|^2$.