

Are correlations constant? Empirical and theoretical results on popular correlation models in finance[☆]



Zeno Adams^{a,*}, Roland Füss^{a,b}, Thorsten Glück^c

^aSwiss Institute of Banking and Finance (s/bf), University of St.Gallen, Unterer Graben 21, 9000 St. Gallen, Switzerland

^bResearch Associate at the Centre for European Economic Research (ZEW), Mannheim, Germany

^cd-fine GmbH, Opernplatz 2, 60313 Frankfurt am Main, Germany

ARTICLE INFO

Article history:

Received 27 June 2016

Accepted 12 July 2017

Available online 14 July 2017

JEL classification:

C12

C52

G01

G11

Keywords:

Change-point tests

Correlation breaks

Dynamic conditional correlation (DCC)

Multivariate GARCH models

Spurious conditional correlation

ABSTRACT

Multivariate GARCH models have been designed as an extension of their univariate counterparts. Such a view is appealing from a modeling perspective but imposes correlation dynamics that are similar to time-varying volatility. In this paper, we argue that correlations are quite different in nature. We demonstrate that the highly unstable and erratic behavior that is typically observed for the correlation among financial assets is to a large extent a statistical artifact. We provide evidence that spurious correlation dynamics occur in response to financial events that are sufficiently large to cause a structural break in the time-series of correlations. A measure for the autocovariance structure of conditional correlations allows us to formally demonstrate that the volatility and the persistence of daily correlations are not primarily driven by financial news but by the level of the underlying true correlation. Our results indicate that a rolling-window sample correlation is often a better choice for empirical applications in finance.

© 2017 Elsevier B.V. All rights reserved.

1. Introduction

Multivariate GARCH models have been designed as extensions of their univariate counterparts. Engle et al. (1984) present an early version as “a bivariate generalization of Engle’s ARCH model”. This view is conceptually appealing and has found widespread use in practice. In this paper, we argue that the nature of dynamic correlations is very different from that of conditional volatilities. While important economic and financial news such as economic activity, interest rate changes, and oil prices affect the volatility of financial assets, the relevance and impact of this news is often similar across firms. As a consequence, volatility is constantly exposed to news and therefore time-varying by nature but correlation changes are only observable after major economic events. For instance, cor-

relations substantially increased for many financial assets following the burst of the Dot-com bubble in 2001 or the default of Lehman Brothers in September 2008 (Ofek and Richardson, 2003; Wied et al., 2012) but correlations are generally insensitive to changes in macroeconomic variables such as interest rates or inflation (King et al., 1994; Karolyi and Stulz, 1996). We demonstrate how current conditional correlation models tend to impose purely artificial dynamics on estimated conditional correlations and show why in empirical applications the estimated parameters governing the dynamics are often statistically significant despite the fact that underlying correlations are constant.

The correlation matrix is the input to many applications in finance and several recent studies seem to believe in the importance of time-varying correlations. For instance, Moskowitz (2003) emphasizes the significance of dynamic conditional correlations during recessions and periods of financial distress. Similarly, Adrian and Brunnermeier (2016) argue that MGARCH models are important for capturing the dynamic evolution of systemic risk. DeMiguel et al. (2009) claim that allowing for time-varying moments could increase the performance of optimal asset allocation. The notion of constant correlations therefore has important implications for financial modeling and practice. Under a constant correlation matrix, international asset portfolios may not have the same

[☆] An earlier version of this paper circulated under the title “Spurious Dynamic Conditional Correlation”. We would like to thank an anonymous reviewer, Carol Alexander (the editor), Chris Brooks, Daniel Buncic, Matthias Fenger, Massimo Guidolin, Michael Massmann, Jan Mutl, Adrian Pagan, David Rapach, and the participants of the SEPS Seminar at the University of St.Gallen for valuable comments and suggestions.

* Corresponding author.

E-mail addresses: zeno.adams@unisg.ch, adams.zeno@gmail.com (Z. Adams), roland.fuess@unisg.ch (R. Füss), info@d-fine.de (T. Glück).




Fig. 1. Daily conditional correlations between S&P 500 and NASDAQ returns.

This figure shows daily DCC correlations between the S&P500 Composite index and the NASDAQ Composite index from 01/01/1990 to 12/31/2014 (6532 obs.). Fluctuations in time-varying correlations are generally high and can change substantially over time. During the period from January 1990 to March 2000 the volatility of daily correlations was 0.077 (122% on an annualized basis) but decreased in the following period (March 2000 to December 2014) to 0.042 (66% annualized).

degree of diversification than comparable portfolios based on dynamic correlations, portfolio optimization could generate different weights, and risk measures may indicate different levels of risk. The aim of our research is not to dismiss dynamic correlation modeling altogether, but to provide a critical perspective on popular models that are routinely used to generate estimates of dynamic asset correlations.

The analysis in this paper is based on Engle's (2002) Dynamic Conditional Correlation (DCC) model. The main advantage of the DCC approach is its parsimonious specification which simplifies interpretation and allows even large asset portfolios to be estimated within seconds. Over the last years, the DCC model has therefore become well-established in both research and practice.¹ In Appendix A of the Online Appendix, we show that our results also hold for other popular MGARCH models, which tend to generate very similar dynamics.² To illustrate its behavior, consider the conditional correlations between the daily returns of the S&P 500 and the NASDAQ index from 1990 to 2014 shown in Fig. 1. Two characteristics that are typical for conditional correlations generated by MGARCH models stand out. First, conditional correlations undergo large swings over a short period of time. In the 1990s, correlations frequently moved within a wide range between 0.52 in July 1993 and 0.94 in November 1997. In the literature, this observation has been sometimes interpreted as evidence that the underlying cor-

relation structure is a highly volatile process (e.g., Pukthuanthong and Roll, 2011; Sadorsky, 2012). Second, the fluctuation in conditional correlations often changes over time. In Fig. 1, correlations are highly volatile during the 1990s but enter a more tranquil period in 2000. During this time, the daily volatility of correlations dropped approximately by half. In this paper, we show that the large fluctuations during the 1990s and the small fluctuations during the 2000s have no fundamental economic cause but are purely artificial results generated by the DCC model. In the following, the discussion and the empirical results of our paper are based on typical bivariate correlations. Simulation results in Pakel et al. (2014) suggest that in the multivariate case, correlations become constant as the number of assets increases.³

We demonstrate that the volatility of estimated conditional correlations $\hat{\rho}$ is a negative function of the underlying true correlation level ρ : The fluctuations in conditional correlations $\hat{\rho}$ are large when the correlation level ρ is close to zero and small when ρ approaches ± 1 .⁴ In Fig. 1, this causes the volatility to decrease drastically as conditional correlations reach values of 0.9 and beyond. In fact, we argue that for financial assets, underlying true correlations ρ are generally constant and that the fluctuations generated by autoregressive-type multivariate GARCH models are spurious. They are caused by infrequent economic disruptions that shift the level of correlations. We recognize that correlations can and do change from time to time. For instance, Longin and Solnik (1995), Bera and Kim (2002), and Forbes and Rigobon (2002) show that correlations among financial assets increase during economic crises and times of financial distress. However, our claim is that these

¹ For instance, the DCC model has been used in value-at-risk estimation (Pérignon and Smith, 2010), the analysis of asset class comovements (You and Daigler, 2010; Heaney and Srikanthakumar, 2012), the implementation of hedging strategies (Chang et al., 2011), and the examination of correlation responses to announcement effects (Brenner et al., 2009), among others.

² Only models that have become accepted in practice and can be applied with reasonable effort and speed are part of our robustness section. This includes MGARCH models with autoregressive covariances such as the corrected DCC model of Aielli (2013), the diagonal VECH model of Bollerslev et al. (1988), or the diagonal BEKK model of Engle and Kroner (1995). It excludes more complex MGARCH specification such as the regime-switching model of Pelletier (2006). For a classification of MGARCH models we refer the reader to Bauwens et al. (2006).

³ We thank an anonymous reviewer for pointing this out.

⁴ In the following, we assume that there exists a true but unobserved correlation ρ . From a statistical viewpoint, our empirical estimates $\hat{\rho}$ are only meaningful if there is a true but unobserved correlation ρ . From an economic perspective, there should be a true correlation coefficient that reflects the way common economic factors lead to comovement between any two financial assets.




Fig. 2. Are daily conditional correlations between S&P 500 and NASDAQ returns constant?

This figure shows the decomposition of daily dynamic correlations into constant correlations that are separated by level shifts. The level shifts are detected by the correlation breakpoint test proposed by Wied et al. (2012). The identified breaks are associated with significant disruptions in financial markets such as the burst of the Dot-com bubble in March 2000 or the failure of Lehman Brothers in September 2008.

level shifts are a very different type of dynamics than the daily autoregressive fluctuations that are suggested by MGARCH models.

To substantiate our claim, we test for breaks in the otherwise constant correlation structure using a recent correlation change-point detection algorithm developed by Galeano and Wied (2014). This algorithm is a repeated application of the change-point test developed by Wied et al. (2012) and is able to identify level shifts that are associated with important financial or economic events. We illustrate this point in Fig. 2, where we repeat the daily DCC correlations from the previous graph but superimpose the constant correlations including their level shifts. The algorithm of Galeano and Wied (2014) suggests that the underlying true correlation was initially constant during the five year period from 1990 to 1994. The change-point tests indicate a shift in the correlation structure in December 1994 which corresponds to the Mexican peso crisis. More correlation breaks followed: the Asian crisis in 1997, the burst of the Dot-com bubble which had its climax in March 2000, and the default of Lehman Brothers in September 2008. The detection algorithm finds no evidence for additional correlation changes within each subsample. We therefore claim that much unlike conditional volatility, the true underlying correlations ρ are likely to be constant.⁵

We can extend this example to other assets. Fig. 3 shows the estimated breaks for the correlation between the S&P 500 and crude oil returns. Two correlation breaks are identified. The first occurs in October 1995 and is fairly small (0.058 in absolute terms). Although it cannot be said with certainty, it is likely that

this break is associated with the Iraq disarmament crisis that involved a series of small events that started in 1995 and escalated in a cruise missile strike that was launched by the U.S. military forces against Iraq air defense targets in September 1996. Since the United States was directly involved in Iraq and the events changed in perception about future Iraq oil production, it is likely to have affected both, the S&P 500 and crude oil prices. The second break is much larger in size (an increase of 0.396) and has recently attracted considerable attention in the empirical literature on commodity markets (e.g. Cheng and Xiong, 2014; Adams and Glück, 2015). At the center of this literature is the comovement between commodities and the stock market that emerged after the bankruptcy of Lehman Brothers. This phenomenon is known as the financialization of commodity markets and has been shown to be responsible for the size and the persistency of the observed comovement (Henderson et al., 2015).

As a third and last example, Fig. 4 shows the correlations between the S&P 500 and Italian government bonds. There are two important events that caused a shift in the level of correlation between these two assets. The first break occurs in 2000 shortly after the introduction of the euro. The euro led to a compression of bond spreads in Italy and other countries in the euro area periphery which could have reduced the correlation between the S&P 500 and Italian government bonds (Hale and Obstfeld, 2016). The second break occurs in 2010 during the sovereign debt crisis when the debt-to-GDP ratio of several European countries increased rapidly. For instance, the debt-to-GDP ratio in Italy increased by almost 20% between 2007 and 2011 (Lane, 2012). In the following years, the yields on Italian government bonds increased to reflect higher risk premia from a deteriorating Italian economy and a higher probability of default (Mody and Sandri, 2012). At the same time, the U.S. stock market recovered from the 2007–2009 financial crisis. As a result, the returns on both, U.S. stocks and Italian government bonds entered a period of positive growth, causing an upward shift in correlations.

⁵ The assumption that daily correlations lie exactly on a straight line between breakpoints may be too strong. For instance, daily trading noise and price fluctuations in S&P500 and NASDAQ stocks is likely to generate small changes in the correlation structure even between breakpoints. However, the Wied et al. (2012) test indicates that those changes are not statistically significant and unlikely to reflect economically relevant changes. Our main arguments in this paper do not depend on the straight line assumption and have the same importance when we allow for correlation noise between breaks.




Fig. 3. Structural breaks in daily correlations between S&P 500 and crude oil returns.

This figure shows two breaks that increased the correlation between U.S. stocks and crude oil returns. The first occurs in the mid-1990s and is associated with the U.S. military involvement in Iraq. The second break coincides with the bankruptcy of Lehman Brothers and marks a strong increase in the comovement between commodities and the stock market. In the empirical literature, this second break is often interpreted as the beginning of a new era of financialization in which the correlation between stocks and commodities are shifted to a permanently higher level.

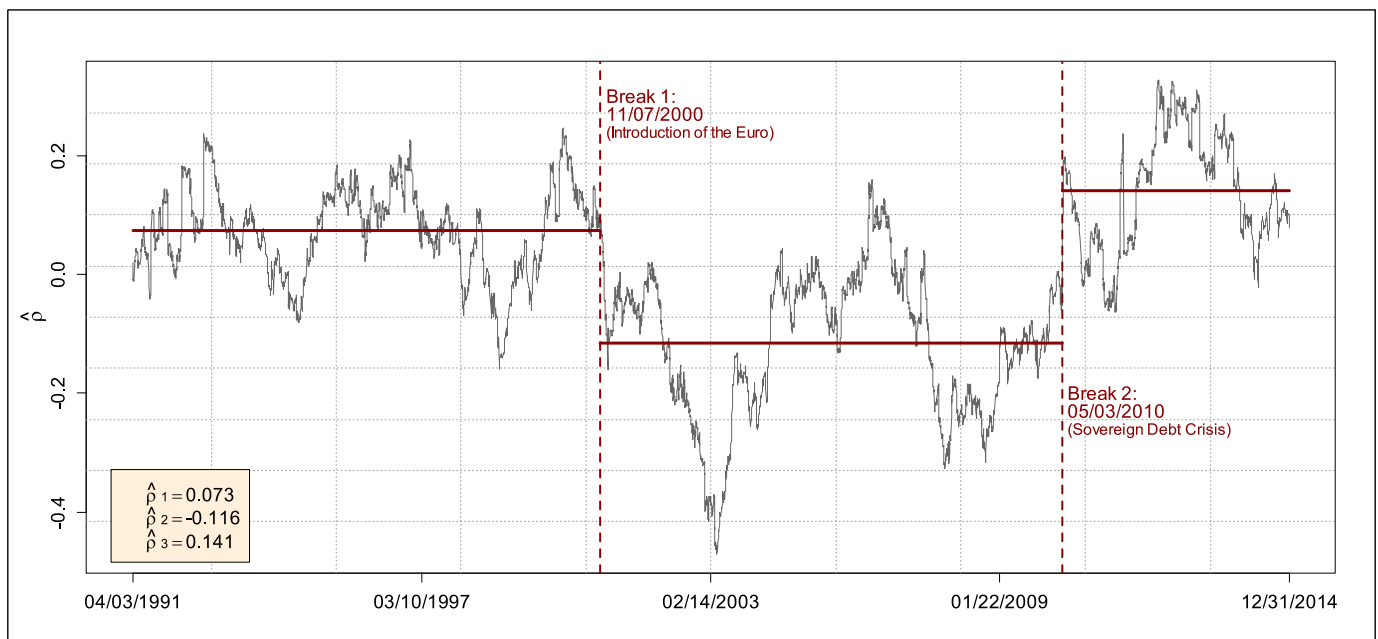


Fig. 4. Structural breaks in daily correlations between S&P 500 and Italian government bonds.

This figure shows two correlation breaks that are associated with the euro. The first break occurs after the introduction of the euro in 1999 and is associated with a compression of bond yields in southern European countries. The second break refers to the sovereign debt crisis when the debt levels of European periphery countries increased.

Our claim that the correlations between financial assets tend to be constant is in line with a number of empirical findings. Tse and Tsui (1999) investigate a number of tests aimed at detecting time-varying correlations through linear dependence in cross products of standardized residuals. They demonstrate that the tests correctly indicate model misspecification when a MGARCH model with constant conditional correlations (Bollerslev, 1990) is fitted to the data

but the true DGP is a MGARCH model with dynamic conditional correlations. However, when applied to empirical data, most studies fail to detect linear dependence (Bollerslev, 1990; Tang, 1995; Tse, 2000). Further evidence is provided by Tse (2000) who derives a Lagrange Multiplier test for constant conditional correlations. He finds this test to have high empirical power if the true DGP is a BEKK or a DCC-type model. Nevertheless, for several assets he fails

to reject the null of constant conditional correlations.⁶ A recent line of literature has proposed MGARCH model extensions that are capable of accommodating breaks between different levels of correlation. For instance, [Silvennoinen and Teräsvirta \(2009\)](#) propose a double smooth transition conditional correlation (DSTCC) model in which the correlations change smoothly between states of constant correlations. The Threshold Conditional Correlation model (TCC) of [Aslanidis and Martinez \(2016\)](#) models abrupt regime changes but has the advantage that it is easier to estimate, even for large numbers of assets. From the empirical evidence produced in this paper, we expect these types of models to gain in importance over the next years.

Given the substantial differences between the volatile DCC correlations in the previous examples and the constant correlations that we claim describes the true underlying correlation structure, an important question is why the DCC parameters that govern the estimated correlation dynamics are often found to be statistically significant in empirical studies. Our empirical results show that statistical significance is much less important than expected. Although statistical significance decreases when we control for correlation breaks, the main difference in correlation dynamics is caused by a change in the size of parameter estimates. These size changes have important consequences for correlation dynamics because the interaction between correlation parameters is highly nonlinear. As a consequence, even small deviations from typical estimates can generate correlations that are either constant or fluctuate randomly at low volatility around a constant value.

A final issue that is important when analyzing correlations is the impact of non-normal return distributions. [Embrechts et al. \(1999\)](#) show that fat-tailed marginal and joint distributions can reduce the attainable range of correlations. While this range is generally between -1 and $+1$, extreme outliers and skewness can push this range towards zero. If the bivariate distribution of standardized returns changes over time to varying degrees of non-normality, the underlying and empirically determined correlations may change as well. This change could be problematic because it can generate a situation in which the true return co-dependence is high but correlations are low. We investigate this phenomenon and its consequences in more detail in Appendix F of this paper.⁷ Our conclusion is that although deviations from non-normality are a salient feature of financial asset returns, the observed deviations may not be sufficiently large to have major distorting effects on our empirical results.

The remainder of this paper is structured as follows. In the next section, we show for several asset classes that correlation breaks among daily returns are a common phenomenon. In [Section 3](#), we take a closer look at the impact of breaks on parameter estimates. When breaks are controlled for, parameter estimates often lie outside the narrow band that produces meaningful correlation dynamics. A theoretical explanation for the results in our paper is explored in [Section 4](#). We derive an expression for the variance and autocovariance of DCC correlations when the true underlying correlations are constant. We show that this expression can be decomposed into a general term that is common to all parametric correlation estimators, and a model-specific term which adopts the unique characteristics of the underlying correlation specification. This allows us to demonstrate how DCC parameters cause artificial

correlation dynamics. In [Section 5](#), we show under which circumstances the historical rolling window sample correlation is preferable to a DCC model. [Section 6](#) summarizes our main results and gives our conclusions.

2. Correlation breaks in daily asset returns

In this paper, we argue that correlations are constant over time, but that financial shocks lead to breaks that shift the level of correlations. This section describes the empirical evidence concerning breaks in the correlation structure of financial assets. Over a 15-year period from 2000 to 2014, the majority of financial assets in our sample experience shocks that significantly shift the level of daily correlations. The presence of such correlation breaks has implications for the estimation of conditional correlation models. The existing literature on univariate volatility models shows that breaks in the volatility dynamics introduce a bias in the estimation which results in inaccurate volatility forecasts ([Hamilton and Susmel, 1994](#); [Hillebrand, 2005](#); [Rapach and Strauss, 2008](#)). However, the findings concerning volatility breaks cannot be simply extended to correlations. We show that correlation dynamics are exposed to model specific factors that are absent in univariate models of volatility.

2.1. The dynamic conditional correlation (DCC) model

Throughout this paper, our emphasis is on [Engle's \(2002\)](#) popular mean-reverting Dynamic Conditional Correlation (DCC) model. The MGARCH family has grown considerably over the last years and a number of more complex models are better at dealing with structural breaks in correlations. For instance, [Mittnik and Paolella \(2000\)](#) propose a weighted maximum likelihood procedure that places less weight on observations in the more distant past. [Pelletier \(2006\)](#) introduces a regime-switching MGARCH model that allows for correlations that are constant within a regime but are different across regimes. The variable governing the changes from one regime to another is a latent unobserved factor. In contrast, the model proposed by [Silvennoinen and Teräsvirta \(2009\)](#) can identify up to two regime switching variables and allows for smooth transitions between breaks. However, the additional flexibility of these more advanced models often comes at the cost of highly sensitive parameter estimates and extensive modeling and forecasting implementation. In this context, the recently proposed Threshold Conditional Correlation model of [Aslanidis and Martinez \(2016\)](#) is promising as it is relatively easy to estimate even for a large number of assets. The results in this paper advocate the use of models that are both practical for large portfolios while at the same time allow for shifts in correlations. We hope that our results can help to promote the kind of models proposed by [Aslanidis and Martinez \(2016\)](#).

For the purpose of our paper, however, we focus on the DCC model which currently enjoys high popularity due to its parsimonious specification and direct interpretation which is analogous to parameters of univariate GARCH models. In other words, we recognize the contribution of the more recent regime switching correlation models but point out that so far they play only a minor role in most empirical settings. For the main arguments in our paper, which are about the dynamic correlations as they are used today, the simple canonical specifications are more relevant. Finally, DCC models generate conditional correlation dynamics which are similar to more complex MGARCH specifications such as BEKK ([Engle, 2002](#); [Engle and Colacito, 2006](#); [Caporin and McAleer, 2012](#)) or the corrected DCC of [Aielli \(2013\)](#). We therefore expect our results to hold also for other members of the MGARCH family.

For our analysis, we abstract from conditional mean effects, i.e. we assume that conditional means are constant. This assumption

⁶ In contrast, comparable studies on the behavior of univariate GARCH models find that GARCH volatility adequately models the true data generating process ([Bollerslev et al., 1992](#)). Overall, the body of empirical literature over the last two decades lends support to the notion that extending the GARCH framework from volatilities to correlations is not as straightforward as previously thought.

⁷ The supplementary internet appendix for this paper is available at http://www.sbf.unisg.ch/en/lehrstuehle/lehrstuhl_fuess/homepage_fuess/publikationen+roland+fuess.

has no serious implications for daily data (Fleming et al., 2001) and is common in the literature (West and Cho, 1995). Furthermore, and for the sake of simplicity, we only consider the bivariate case. This restriction does not affect our results as DCC model parameters have the same impact on each component of the conditional correlation matrix. Hence, all our results apply to higher dimensions as well. The bivariate DCC model uses the specification

$$\rho_t = f(\mathbf{q}_t) = \frac{q_{12,t}}{\sqrt{q_{11,t} \cdot q_{22,t}}}, \quad (1)$$

where the elements of $\mathbf{q}_t = (q_{11,t}, q_{22,t}, q_{12,t})$ are functions of the standardized residuals $\{(e_{1,k}, e_{2,k})' : k < t\}$. The residuals $e_{i,t}$ are obtained as the GARCH residuals from the return equation $r_{i,t} = \sigma_{i,t} e_{i,t}$ and are typically associated with financial shocks or news (Engle and Ng, 1993). The q_t are estimated using the expression:

$$\begin{aligned} \hat{q}_{i,j,t} &= \frac{1-a-b}{1-b} \vartheta_{i,j} + a \sum_{s=1}^{\infty} b^{s-1} e_{i,t-s} e_{j,t-s} \\ &= (1-a-b) \vartheta_{i,j} + a e_{i,t-1} e_{j,t-1} + b \hat{q}_{i,j,t-1} \end{aligned} \quad (2)$$

with $a > 0$, $b > 0$, and $a + b < 1$. Intuitively, $q_{11,t}$ and $q_{22,t}$ can be regarded as auxiliary estimates of $E(e_{1,t}^2 | \mathcal{F}_{t-1})$ and $E(e_{2,t}^2 | \mathcal{F}_{t-1})$, respectively, where \mathcal{F}_{t-1} denotes the information set available at time $t-1$. Similarly, $q_{12,t}$ is an auxiliary estimate of $E(e_{1,t} e_{2,t} | \mathcal{F}_{t-1})$. In line with this interpretation, Eq. (1) simply states rescaling in such a way that $\hat{\rho}_t$ satisfies $-1 \leq \hat{\rho}_t \leq 1$. If the unconditional correlation between $e_{i,t}$ and $e_{j,t}$ is the same for all t , then a natural choice for the constant $\vartheta_{i,j}$ is $E(e_{i,t} e_{j,t})$. Although this definition of the constant does not assure that $E(\hat{\rho}_t) = E(e_{i,t} e_{j,t})$, most empirical applications are based on this specification. For this reason and because setting $\vartheta_{i,j} = E(e_{i,t} e_{j,t})$ considerably simplifies the analysis, we hence refrain from adjustments suggested in Aielli (2013) and refer to Caporin and McAleer (2008) for a detailed discussion. The parameter a models the sensitivity of q_t to the arrival of news $e_{i,t}$. If a is close to zero, correlations dynamics are insensitive to shocks and approximate a straight line. The parameter a therefore plays a special role in our paper and is labeled “news parameter”. In contrast, the parameter b measures the persistence in correlations and is referred to as the “decay parameter”. A low b value generates correlations that fluctuate randomly at low volatility around a straight line.⁸

2.2. A correlation breakpoint test

To identify and locate change points in the correlation structure of financial assets we implement a simple and effective algorithm proposed by Galeano and Wied (2014) and Wied et al. (2012). Similar to recent correlation regime switching models (e.g., Silvennoinen and Teräsvirta, 2009; Aslanidis and Martinez, 2016) the methodology does not require possible break dates to be specified in advance but uses an algorithm for sequential breakpoint detection.⁹ This algorithm involves the following steps: Consider a sample of T observations of the return vector $(r_{1,t}, r_{2,t})'$. Let ρ_t denote the true but unknown unconditional correlation between $r_{1,t}$ and $r_{2,t}$ at time t . The algorithm tests the null hypothesis of constant correlations against the alternative hypothesis of a change-point t^c , i.e.

$$H_0 : \rho_t = \rho \quad \text{for all } t \in \{1, \dots, T\} \quad (3)$$

⁸ An alternative measure of persistence that is sometimes used in the univariate GARCH literature is the sum of a and b (Hillebrand, 2005). Although the effectiveness of b depends on a reasonable value for a (somewhere between 0.02 and 0.06) the actual persistence is governed by the size of b . We therefore rely on b as a measure of persistence throughout the paper.

⁹ An alternative correlation break point methodology is the test proposed by Andreou and Ghysels (2002, 2003). The results in our paper do not change significantly under this alternative test. To conserve space, we do not report the results here but instead focus on the test of Wied et al. (2012).

versus

$$H_1 : \exists t^c \in \{1, \dots, T-1\} \quad \text{such that } \rho_{t^c} \neq \rho_{t^c+1}. \quad (4)$$

The procedure is based on the model-free fluctuation-type test (“WKD test” henceforth) originally proposed by Wied et al. (2012). The test statistic is defined as

$$Q_T := \hat{D} \max_{2 \leq t \leq T} \frac{t}{\sqrt{T}} |\hat{\rho}_t - \hat{\rho}_T|, \quad (5)$$

where $\hat{\rho}_t$ is the sample correlation over the period 1 to t . The purpose of the scalar coefficient \hat{D} is to rescale the volatility of $\hat{\rho}_t$ which tends to be higher at the beginning of the sample when only a few observations are available. The coefficient \hat{D} is described in more detail in Appendix B of the Online Appendix. Under the null hypothesis and several reasonable moment and dependency restrictions, the test statistic Q_T is asymptotically Kolmogorov distributed (Wied et al., 2012, Theorem 1).¹⁰ If Q_T stays below the upper critical value the null hypothesis of constant correlation cannot be rejected and the algorithm stops. Otherwise, H_0 is rejected and the correlation sample contains at least one change-point t^c . The estimator for the single change-point is defined as

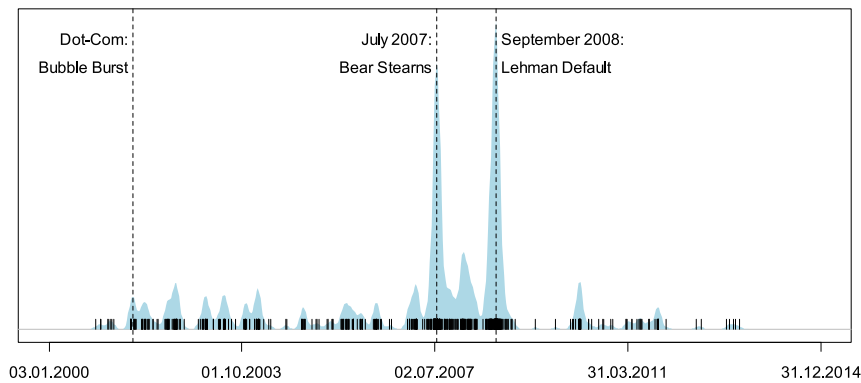
$$t^c = \arg \max_t \hat{D} \frac{t}{\sqrt{T}} |\hat{\rho}_t - \hat{\rho}_T|. \quad (6)$$

To identify further change-points, the sample is split into the two subsamples $[1, \dots, \hat{t}^c]$ and $[\hat{t}^c + 1, \dots, T]$. These subsamples are then both tested individually. This procedure is repeated until no further change-points are detected. Galeano and Wied (2014) demonstrate that the presence of multiple change-points can affect the test’s efficiency in identifying the true number of change points. The last step of the algorithm therefore consists of a refining process in which the vector of the n detected change-points $\tau = [\hat{t}_1^c, \dots, \hat{t}_n^c]$, sorted in ascending date order $\hat{t}_1^c \leq \dots \leq \hat{t}_n^c$, is verified in subsamples containing only a single change point. For the implementation of the refining process we define the first observation of the sample as $\hat{t}_0^c = 0$, the last observation as $\hat{t}_{n+1}^c = T$, and form the subsamples $[\hat{t}_{i-1}^c + 1, \dots, \hat{t}_i^c]$ for $i = 1, \dots, n$. Each subsample starts at the first observation following the previous change-point \hat{t}_{i-1}^c , includes change point \hat{t}_i^c , but ends just before the next change-point \hat{t}_{i+1}^c . These subsamples are tested individually. If the null hypothesis is not rejected the change-point contained in the subsample is removed from τ .


We give a brief example to illustrate the test procedure. We test for a break in the daily correlation between the returns of the FTSE100 index and the returns of the Dutch AEX index. If we test the full sample from 01/03/2000 to 11/28/2014 we detect one significant change-point at 01/17/2008. In the next step, we split the data in the two subsamples [01/03/2000–1/17/2008] and [01/18/2008–11/28/2014]. We test the two subsamples individually and find another change-point in the first subsample at 08/31/2001. The test statistic in the second subsample is insignificant. The presence of the second change-point interferes with the test statistic for the first change-point that was detected using the entire sample. The last step therefore involves a refining process in which we test the subsamples [01/03/2000–01/17/2008] and [09/01/2001–11/28/2014]. The test statistics for both subsamples remains significant and confirms the presence of both change-points: $\tau = [08/31/2001, 01/17/2008]$. Galeano and Wied (2014) demonstrate that this procedure detects the correct number of correlation change-points.

We apply the WKD test to the daily returns of 40 assets over the period 01/03/2000 to 12/31/2014 (3914 obs.). Our data covers

¹⁰ See assumptions A.1–A.5 in Wied et al. (2012). In particular, it is assumed that $\{(r_{1,t}, r_{2,t})'\}$ is near-epoch dependent. For an extensive discussion see Davidson (1994, Ch. 17).

Panel A: Distribution of Correlation Breaks Over Time**Panel B: Number of Correlation Breaks (2000 – 2014)**

0	1	2	3	4	5
201	420	84	47	16	6
(26%)	(54%)	(10%)	(6%)	(2%)	(1%)

Panel C: Change in Average Correlation $\hat{\rho}$ Following a Break**Fig. 5.** Correlation breaks: descriptive statistics.

This figure shows the distribution of correlation breaks over time (Panel A), the frequency of breaks (Panel B), and the average change in correlations following a break (Panel C). Breaks are estimated over 40 assets including stocks, bonds, commodities, and currencies for a total of $(40 \times 39)/2 = 780$ correlation pairs. The WKD test detects 775 breaks in daily correlations over the period 2000 to 2014.

the asset classes stocks, bonds, commodities, and currencies. Each asset class is represented by 10 major indices or currency pairs. Appendix C in the Online Appendix lists the constituents in detail. We obtain $1/2 \cdot n(n-1) = 780$ correlation time series to be tested for breaks. Panel A in Fig. 5 shows the distribution of breaks over time. Intuitively, correlation breaks should cluster around dates that are associated with important financial shocks. Panel A indicates that this is in fact the case. The two events that appear to have influenced correlations most are the failure of two Bear Stearns funds in July 2007 and the bankruptcy of Lehman Brothers in September 2008. If we look at the table of correlation breaks in Panel B we see that 420 return pairs, or more than half of the assets in our data experienced exactly one correlation break. More than two correlation breaks is much less common. In 26% of cases, the WKD test did not identify a significant change in correlations.¹¹ Finally, Panel C shows how correlations change after a break. Since small changes are likely to be statistically insignificant the distri-

bution shows a distinct bimodal shape. Financial crises tend to increase the comovement among assets which can explain the higher positive mode. In fact, the majority, or 61% of correlation changes are positive. The largest negative drop in average correlations is -0.68 . The largest positive jump is as high as 0.58 . The average positive or negative correlation change is around 0.18 . From our findings in Fig. 5 we conclude that correlation breaks in financial assets is a common phenomenon and that the change in correlations following a break is often large. In the following section, we investigate the consequences of breaks on DCC parameter estimates which govern the dynamics of conditional correlations.

3. The impact of correlation breaks on DCC parameter estimates

Over the 15 year period from 2000 to 2014, the majority of financial assets experienced at least one correlation break. If correlation breaks are a prevalent characteristic of financial assets, the question is how this affects the parameters that govern the correlation dynamics. To answer this question, we use the $n = 40$ assets from the previous section and obtain $1/2 \cdot n(n-1) = 780$ cor-

¹¹ This appears to be particularly the case for assets that had very low correlations to begin with.

relation time series. First, we collect the DCC parameter estimates over the full sample from 03/01/2000 to 12/31/2014 (3914 obs.). In a second step, we re-estimate the parameters running the model over the subsamples between correlations breaks.¹² We obtain 780 DCC parameter pairs for the full samples and 1596 parameter pairs for the subsamples. In the following, the full sample estimates serve as our control group while the subsample estimates are our treatment group. To allow for a direct comparison between the two groups we need to ensure that the full sample correlations contain at least one correlation break. In addition, we exclude correlations that are already constant in full samples. This occurred in a number of cases in which the news parameter was close to zero.¹³ We discuss this important situation in more detail in Section 4. Finally, we remove subsamples that contain less than 500 observations to ensure that our results are not driven by small sample windows. Our final sample consists of 355 full sample correlations and 802 subsample correlations.

Panel A of Fig. 6 shows the descriptive statistics of DCC parameters and compares full samples to subsamples. We find an average value for the news parameter a of 0.021 and for the decay parameter b of 0.97 which are in line with findings from other studies. For instance, Engle and Sheppard (2001) investigate the daily returns of the DJIA index and report 0.01 and 0.96 for a and b , respectively. Engle and Colacito (2006) investigate the daily returns of the S&P 500 and the 10-year U.S. bond futures contract and report 0.022 and 0.973 for a and b . Measured at the 95% confidence level, our parameter estimates are statistically significant in 91% of all cases.

When we compare these numbers to the estimates found in subsamples the average coefficients appear to be quite similar. The average news parameter is close to the full sample estimate while the decay parameter estimate is somewhat lower. This observation is in line with Rapach and Strauss (2008) and Hillebrand (2005) who find that the persistence parameter β for univariate GARCH models decreases after accounting for structural breaks in volatility. The main change in subsamples parameters, however, does not take place in the average estimate but in its distribution. For instance, subsample estimates for the news parameter now range between 0 and 0.244 while estimate for the decay parameter range between 0 and 0.997. Panel B of Fig. 6 illustrates the impact of correlation breaks on the parameter distribution for different asset classes in our sample. The distribution of news parameters shows a number of positive outliers, in particular in cases when the correlations are measured among stocks, bonds, or between mixed assets types. This finding is even stronger when we look at the change in distribution of the decay parameter. A significant part of the distribution now covers parameters below 0.5 which leads to a low persistence in the correlation dynamics. To conclude, estimating DCC models in subsamples that contain no breaks has subtle but important consequences for DCC parameter estimates. However, it is unclear how the changes in DCC parameters affect correlation dynamics. We will explore this issue in the following section.

4. The impact of correlation breaks on correlation dynamics

In this section, we show how breaks in the correlation structure of financial assets affect the dynamics of correlation estimators. We derive an expression for the autocovariance function of conditional correlations. This expression shows how the autocovariances can be separated into a general variance part and a model

part which explicitly depends on the parameterization of the correlation process. In this context, we compare the results from the DCC model to two simple correlation estimators that are popular in practice: the fixed parameter exponentially weighted moving average (EWMA) estimator and the historical sample correlation estimated in a rolling window.¹⁴ In a number of cases, these simple alternatives will perform better than a MGARCH-type correlation estimator.

4.1. Dynamic correlation estimators

Our first estimator is the *rolling window (RW) sample correlation*. The (historical) sample correlation of asset returns is a simple and popular estimator in applied finance. However, it is not just an ad hoc way to measure dynamic correlations. As shown in Foster and Nelson (1996) there are several DGPs for which an appropriately specified rolling window estimator is optimal. We follow the literature and use this estimator as a benchmark to be tested against the performance of more sophisticated MGARCH models (Engle, 2002). Again, we can facilitate comparison among estimators by expressing the RW correlation in terms of $\hat{q}_{i,j,t}$, which in the case of the rolling window estimator is a function of n equally weighted observations ranging from $t - n$ to $t - 1$:

$$\hat{q}_{i,j,t} = n^{-1} \sum_{s=1}^n e_{i,t-s} e_{j,t-s}. \quad (7)$$

Our second correlation estimator is the fixed parameter *Exponentially Weighted Moving Average (EWMA)* estimator. The implementation of this estimator is just slightly more elaborate than a rolling window correlation. The dynamics generated by the EWMA model are very similar to MGARCH models but the parameters are given rather than being estimated.¹⁵ RiskMetrics suggests modeling the dynamics in the daily asset return covariance using a persistence parameter $\lambda = 0.94$ (Longerstaey and More, 1995). The response to shocks is measured by the remaining $(1 - \lambda) = 0.06$

$$\hat{q}_{i,j,t} = (1 - \lambda) \sum_{s=1}^{\infty} \lambda^{s-1} e_{i,t-s} e_{j,t-s} = (1 - \lambda) e_{i,t-1} e_{j,t-1} + \lambda \hat{q}_{i,j,t-1}. \quad (8)$$

The importance of the standardized shocks from both assets $e_{i,t-s} e_{j,t-s}$ decreases exponentially, thereby emphasizing the information provided by current observations relative to past observations. For the analysis in our paper, the EWMA model is useful in two ways. First, the effort of implementing the EWMA estimator is somewhere between a rolling window estimator and the more sophisticated MGARCH models. Second, since the persistence parameter λ and the shock sensitivity $1 - \lambda$ sum up to one, the resulting correlations are non-stationary, so that shocks can generate permanent level shifts in dynamic correlations (Engle, 2002). This feature is important since economic disruptions are reflected in many financial time series. Although the EWMA model suffers from similar shortcomings as the DCC model, it should in principle be particularly suited for this purpose.

Our last estimator is the *Dynamic Conditional Correlations (DCC)* estimator that was already presented in Section 2.1. Because of the constant in the function of $q_{ij,t}$, Engle's (2002) mean-reverting DCC specification can be considered as an extension of the EWMA model. For the sake of completeness, we repeat the expression for

¹² We exclude five observations before and after the break to remove possible transition effects that occur around the break date.

¹³ We impose the parameter restrictions $0.01 < a < 0.06$ and $0.8 < b < 0.99$ which will generate a typical dynamic correlation behavior.

¹⁴ The EWMA estimator uses fixed parameters proposed by the RiskMetrics group and is therefore sometimes called the "RiskMetrics model" (Rapach and Strauss, 2008).

¹⁵ When the parameters are estimated the EWMA becomes the Integrated Dynamic Conditional Correlations (IDCC) estimator (Engle 2002).

Panel A: Descriptive Statistics of DCC Parameters					
	Mean	SD	Min	Max	Significant
Full Sample					
News Par.	0.021	0.010	0.010	0.059	91%
Decay Par.	0.970	0.018	0.847	0.988	100%
Subsample					
News Par.	0.027	0.023	0.000	0.244	58%
Decay Par.	0.836	0.276	0.000	0.997	87%

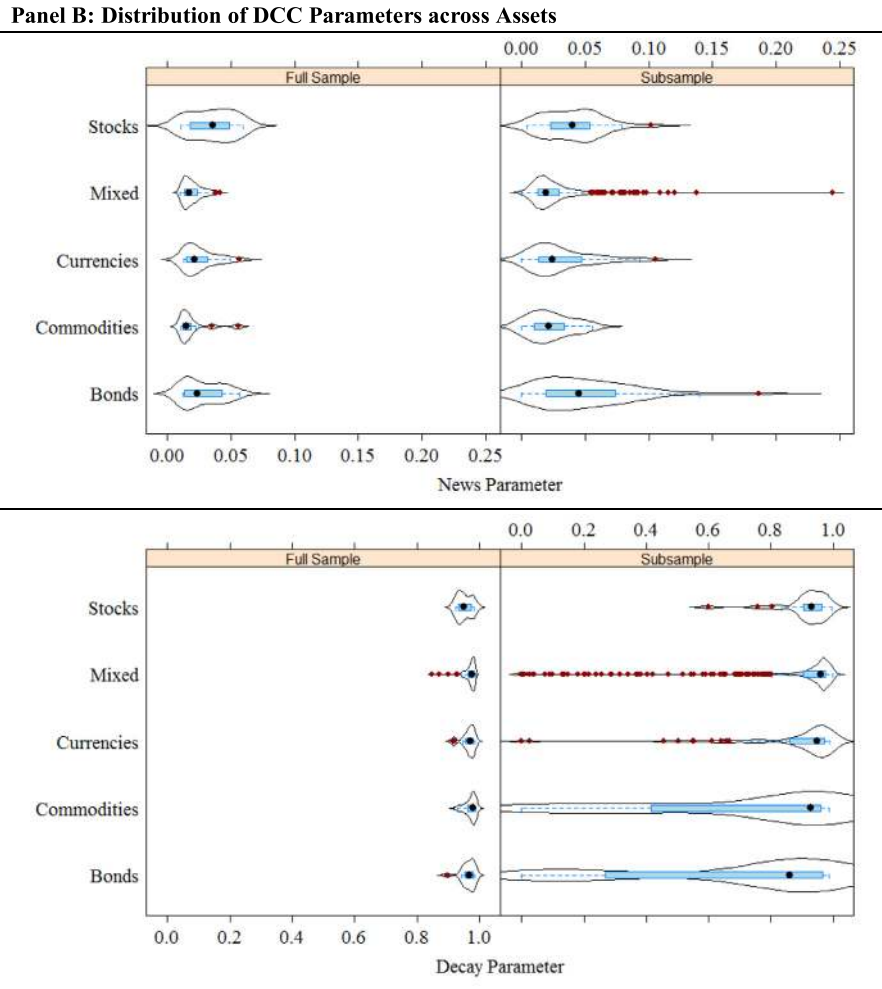


Fig. 6. DCC parameters in full samples and subsamples.

This figure compares the behavior of DCC parameters in full samples and subsamples. The full sample consists of correlations that experience at least one break from 2000 to 2014 (74% of our sample) and have DCC parameters within typical ranges: $0.01 < a < 0.06$ and $0.8 < b < 0.99$ (62% of our sample or 355 correlation series). The parameters in both samples are based on the same asset pairs to allow a direct comparison.

$\hat{q}_{i,j,t}$ here:

$$\hat{q}_{i,j,t} = \frac{1-a-b}{1-b} \vartheta_{i,j} + a \sum_{s=1}^{\infty} b^{s-1} e_{i,t-s} e_{j,t-s} = (1-a-b) \vartheta_{i,j} + a e_{i,t-1} e_{j,t-1} + b \hat{q}_{i,j,t-1}. \tag{9}$$

We will now investigate how changes in the correlation parameters govern the dynamics of correlation estimators.

4.2. Anatomy of conditional correlation dynamics

In this section, we show that the volatile pattern that is typically observed when applying conditional correlation measures is to a large extent artificial. We demonstrate that correlation dynamics depend on the specification of the correlation model. The

impact of the model in turn is amplified by the level of the underlying correlation structure. The volatility of the estimated dynamic correlation $\hat{\rho}_t$ is highest when the true underlying correlation ρ is zero and diminishes as ρ approaches ± 1 . The key finding is that estimated correlations contain a fluctuation component that is unrelated to the correlation parameters. In the following we will formally demonstrate this effect. We start by considering a DGP that produces constant correlations. In terms of \mathbf{q}_t as in Eq. (8) above we define: $\rho_t = f(\mathbf{q}_t) = q_{12,t} / \sqrt{q_{11,t} q_{22,t}}$, where $\mathbf{q}_t = (1, 1, \rho)$. Since the true correlation is constant, it follows that $\rho = E(e_{1,t} e_{2,t} | \mathcal{F}_{t-1}) = E(e_{1,t} e_{2,t})$ and the variance and all autocorrelations of conditional correlations are zero.

To focus on the correlation effect, we assume that we have a correctly specified model for the conditional return volatility, so

that the standardized residuals e_t are homoscedastic and normally distributed:

$$(e_{1,t}, e_{2,t})' \sim i.i.d.N\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}\right). \quad (10)$$

In the following, we will simplify the analysis and use the first-order Taylor expansion of $f(\hat{\mathbf{q}}_t)$ around $E(\mathbf{q}_t) = \varphi = (1, 1, \rho)$ to approximate $Cov(\hat{\rho}_t, \hat{\rho}_{t-s})$. A Monte Carlo simulation in Appendix D of the Online Appendix shows that a reasonable choice of model parameters approximates the exact analytical expression for $Cov(\hat{\rho}_t, \hat{\rho}_{t-s})$ quite well.¹⁶ The Taylor series expansion of $\hat{\rho}_t^*$ can be expressed as

$$\hat{\rho}_t^* = f(\varphi) + \frac{\partial f(\varphi)}{\partial f\varphi'}(\hat{\mathbf{q}}_t - \varphi). \quad (11)$$

We will therefore approximate $Cov(\hat{\rho}_t, \hat{\rho}_{t-s})$ by

$$Cov(\hat{\rho}_t^*, \hat{\rho}_{t-s}^*) = \frac{\partial f(\varphi)}{\partial f\varphi'} Cov(\hat{\mathbf{q}}_t, \hat{\mathbf{q}}_{t-s}) \frac{\partial f(\varphi)}{\partial f\varphi'}. \quad (12)$$

Proposition 1. Let $\{(e_{1,t}, e_{2,t})'\}$ be a bivariate i.i.d. process as defined in (10) with $|\rho| < 1$ and $s \geq 0$, and define

$$\mathbf{\Omega} = 2 \begin{pmatrix} 1 & \rho^2 & \rho \\ \rho^2 & 1 & \rho \\ \rho & \rho & \frac{1}{2}(1 + \rho^2) \end{pmatrix}.$$

If $\hat{\mathbf{q}}_t$ follows the definition in (9) then

$$(i) \quad Cov(\hat{\mathbf{q}}_t, \hat{\mathbf{q}}_{t-s}) = \gamma^{DCC}(a, b, s)\mathbf{\Omega}, \text{ where } \gamma^{DCC}(a, b, s) = \frac{b^s a^2}{1-b^2}.$$

If $\hat{\mathbf{q}}_t$ follows the definition in (8) then

$$(ii) \quad Cov(\hat{\mathbf{q}}_t, \hat{\mathbf{q}}_{t-s}) = \gamma^{EWMA}(\lambda, s)\mathbf{\Omega}, \text{ where } \gamma^{EWMA}(\lambda, s) = \frac{\lambda^s(1-\lambda)^2}{1-\lambda^2}.$$

If $\hat{\mathbf{q}}_t$ follows the definition in (7) then

$$(iii) \quad Cov(\hat{\mathbf{q}}_t, \hat{\mathbf{q}}_{t-s}) = \gamma^{RW}(n, s)\mathbf{\Omega}, \text{ where } \gamma^{RW}(n, s) = \frac{n-s}{n^2} \text{ for } s < n \text{ and } \gamma^{RW}(n, s) = 0 \text{ for } s \geq n.$$

Proposition 1 implies that $Cov(\hat{\rho}_t^*, \hat{\rho}_{t-s}^*)$ contains a true unconditional correlation component $\frac{\partial f(\varphi)}{\partial f\varphi'} \mathbf{\Omega} \frac{\partial f(\varphi)}{\partial f\varphi'}$, and a model-specific multiplier γ^{DCC} , γ^{EWMA} , and γ^{RW} . It is important to note that $\mathbf{\Omega}$ is unrelated to any model parameters. Fluctuations over time are therefore an inherent part of correlation dynamics (see also the proof of Proposition 1 in Appendix E of the Online Appendix).

Proposition 1 allows us to express $Cov(\hat{\rho}_t^*, \hat{\rho}_{t-s}^*)$ in a single number. Let $\kappa(\rho) = (1 - \rho^2)^2$. For the case that $\hat{\mathbf{q}}_t$ follows the DCC model we obtain

$$Cov(\hat{\rho}_t^*, \hat{\rho}_{t-s}^*) = \gamma^{DCC}(a, b, s)\kappa(\rho).$$

For the case that $\hat{\mathbf{q}}_t$ follows the EWMA model we obtain

$$Cov(\hat{\rho}_t^*, \hat{\rho}_{t-s}^*) = \gamma^{EWMA}(\lambda, s)\kappa(\rho).$$

Finally, for the case that $\hat{\mathbf{q}}_t$ follows the RW model we obtain

$$Cov(\hat{\rho}_t^*, \hat{\rho}_{t-s}^*) = \gamma^{RW}(n, s)\kappa(\rho).$$

For the discussion of correlation dynamics we focus on these expressions of $Cov(\hat{\rho}_t^*, \hat{\rho}_{t-s}^*)$.

4.3. The impact on correlation dynamics

We have demonstrated that $Cov(\hat{\rho}_t^*, \hat{\rho}_{t-s}^*)$ has a model-independent component $\kappa(\rho)$ that is fully determined by the size of the true underlying correlation ρ , and a model multiplier γ . In other words, a dynamic correlation estimator $\hat{\rho}_t$ generates spurious dynamics even when the true underlying correlation ρ is constant.¹⁷ With $|\rho| < 1$ and $\gamma > 0$, setting the lag s to zero results

¹⁶ For further examples, see Kwan (2008) and the references therein. The internet appendix for this paper is available at http://www.sbf.unisg.ch/en/lehrstuehle/lehrstuhl_fuess/homepage_fuess/publikationen+roland+fuess.

¹⁷ Technically, $\hat{\rho}_t$ will be zero if $\gamma = 0$. We can ignore this situation since this describes the trivial case when all parameters driving the correlation dynamics are zero. Note that a positive rolling window multiplier requires $s < n$.

in the variance of $\hat{\rho}_t^*$: $Cov(\hat{\rho}_t^*, \hat{\rho}_t^*) = Var(\hat{\rho}_t^*) > 0$. The definition of $\kappa(\rho) = (1 - \rho^2)^2$ shows that, as long as returns are not perfectly correlated, $\kappa(\rho)$ is always positive and has a global maximum at $\rho = 0$. As a consequence, the estimated correlation $\hat{\rho}_t$ contains daily fluctuations even when the true correlation does not. In addition, $\hat{\rho}_t$ depends on lagged values when $Cov(\hat{\rho}_t^*, \hat{\rho}_{t-s}^*)$ is significantly positive for lags $s > 0$. A high persistence can give the pattern in $\hat{\rho}_t$ a spurious dynamic behavior that some researchers have interpreted as cycles (e.g., Cai et al., 2009; Pukthuanthong and Roll, 2011).¹⁸

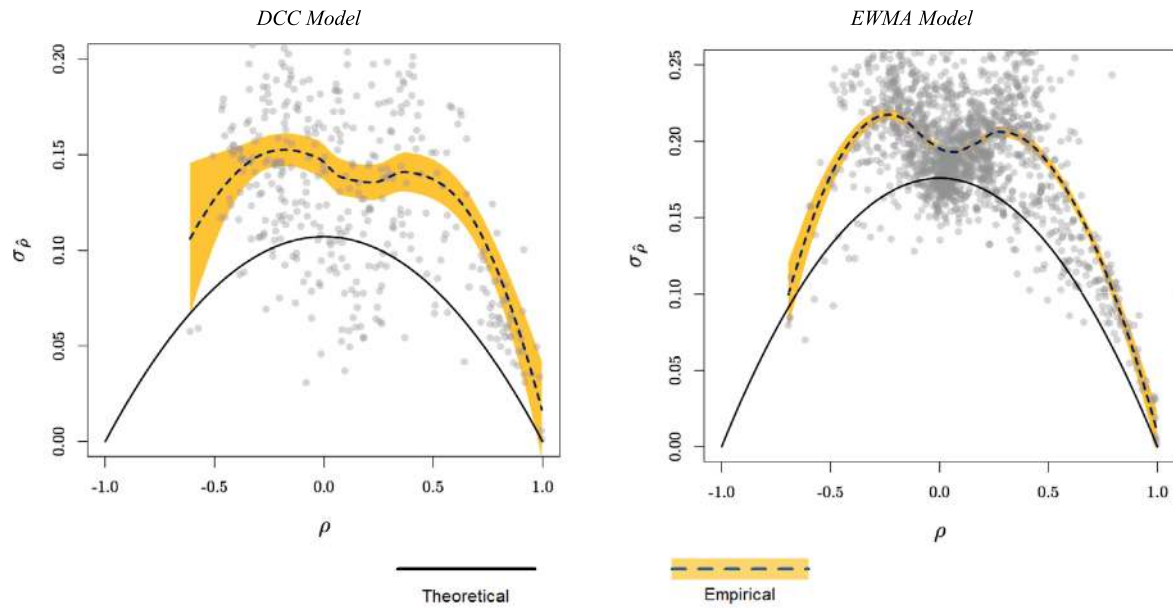
Given the correlation parameters (a and b in the DCC model, λ in the EWMA model, and n in the rolling window model) the relationship between the volatility of the correlation estimate $\sigma_{\hat{\rho}} \approx \sqrt{Var(\hat{\rho}_t^*)} = \sqrt{\gamma\kappa(\rho)}$ and the level of the true underlying correlation ρ is described by an inverse parabolic relationship. A high correlation level such as ± 0.95 generates stable correlation dynamics irrespective of the underlying model specification. If the level of correlation decreases from 0.95 to 0.50 the volatility of $\hat{\rho}_t$ increases by a factor of 6. If the underlying correlation decreases further to 0, the volatility of $\hat{\rho}_t$ increases by another 33%. Again, this increase in volatility occurs for all model specifications and parameter choices. In empirical applications, the pronounced fluctuations in correlation dynamics are often interpreted as the result of incoming financial news (Christiansen, 2000; Capiello et al., 2006). A simple and more likely explanation is that assets share a low level of correlation. We can test whether this relationship can also be found in empirical data. To obtain an estimate of the volatility in the correlation measure, $\hat{\sigma}_{\hat{\rho}}$, we estimate dynamic correlations $\hat{\rho}_t$ for our 40 assets and take the sample standard deviation. We can obtain a measure of ρ by taking the average over the time series of $\hat{\rho}$.¹⁹ Panel A of Fig. 7 plots $\hat{\sigma}_{\hat{\rho}}$ on the y-axis against the level of correlation on the x-axis. The solid line shows the theoretical relationship of $\hat{\sigma}_{\hat{\rho}}$ as derived in the previous section emphasizing that $\hat{\sigma}_{\hat{\rho}}$ is an inverse parabolic function of ρ . For instance, $\hat{\sigma}_{\hat{\rho}}$ in the DCC model is generated by $\hat{\sigma}_{\hat{\rho}} = \sqrt{\gamma^{DCC} \cdot \kappa(\rho)} = \sqrt{0.03^2 / (1 - 0.96^2) \cdot (1 - \rho^2)^2}$, where we have used typical DCC parameters (0.03 for the news parameter and 0.96 for the decay parameter). The hump-shaped relationship is shown in the left graph of Panel A. For comparison, the dashed line with 95% confidence bands shows an estimated nonparametric relationship based on the actual data. The empirical relationship matches the theoretical one, lending support to the notion that dynamic correlations are likely to be more imprecise when the underlying correlation is close to zero.²⁰ The EWMA model in the right graph of Panel A verifies this observation. Although we lack data for extreme negative correlations, we find strong evidence that the relationship holds for large positive correlations. An increasing variation in the observations around $\rho = 0$ suggests that the link between $\hat{\sigma}_{\hat{\rho}}$ and ρ is less clear when assets are uncorrelated.

¹⁸ We note several similarities between our findings and those on rolling windows and ARMA-type processes provided elsewhere (e.g., Lütkepohl, 2006). However, except for some simulation results reported in Aielli (2013), we are not aware of any paper considering the autocovariance structure of conditional correlations.

¹⁹ The presence of correlation breaks is likely to change the model multiplier γ . To separate this effect from the hump shaped function of the model-independent component $\kappa(\rho)$, we base our dynamic correlation estimates on subsamples.

²⁰ For the empirical implementation, we had to assume that the average correlation generated by the DCC model is a good approximation of the underlying true correlation level ρ . Embrechts et al. (1999) show that the attainable range of ρ is reduced when the bivariate distribution of returns deviates strongly from normality. This can also affect the estimated form of the hump-shaped relationship. In Appendix F of the Online Appendix, we show that the relationship looks very similar when estimated with the distribution free Kendall's τ . This suggests that although deviations from non-normality are common, they are unlikely to have strong distorting effects on our results.

Panel A: Variation in $\hat{\rho}$ as a Function of the Underlying ρ : Theory and Empirical Evidence



Panel B: Simulated Relationship and the Connection to Volatility: A Comparison

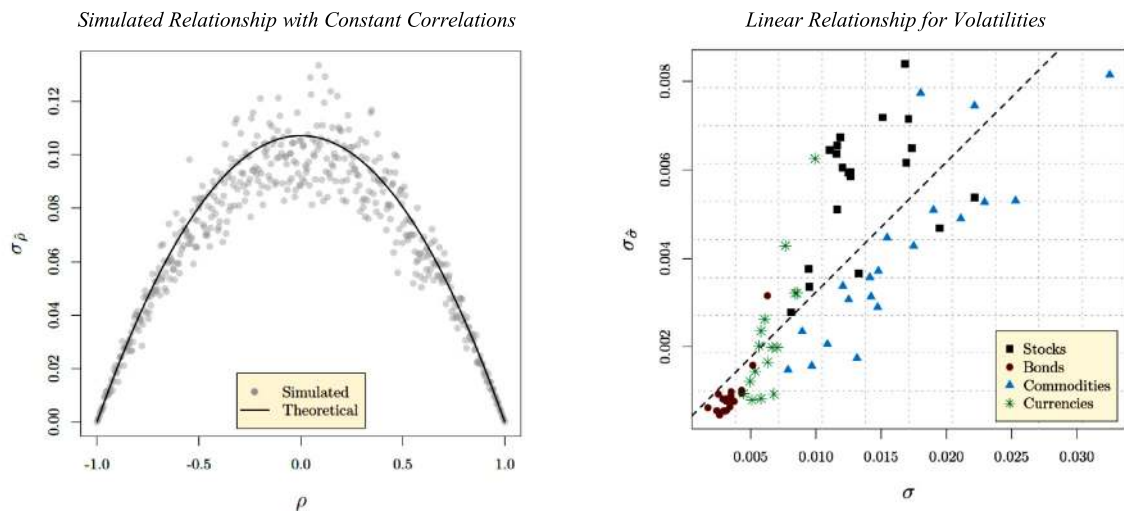


Fig. 7. The parabolic relationship between the level and volatility of dynamic correlations.

This figure shows how typical estimators of dynamic correlations, $\hat{\rho}$, are influenced by the levels of the true but unobserved correlation ρ . Panel A compares the theoretical behavior, indicated by the inverse parabolic relationship (solid line) with the actual behavior found in the data (dashed line). The theoretical relationship for the DCC model is based on a news parameter of 0.03 and a decay parameter of 0.96. The theoretical relationship for the EWMA model is based on $\lambda = 0.94$ and $1 - \lambda = 0.06$. The data consists of various correlation pairs from common indices representing stocks, bonds, commodities, and currencies. The points in the left graph show the means and standard deviations of 404 DCC correlations from this group of assets. The right graph shows the estimates of the EWMA correlations and consists of 1613 observations. The lower number of observations for the DCC model results from excluding correlations with news parameter estimates of less than 0.02 and decay parameter estimates of less than 0.8. Panel B gives an example of the relationship in a controlled simulated environment and shows the case for univariate volatility models. The left graph shows the simulated relationship between variation and level of ρ when the true underlying ρ is constant over time. The right graph shows that the pronounced relationship between the level and volatility of ρ is unique to dynamic correlation models and does not extend to univariate GARCH models.

The theoretical relationship is derived under our assumption that true correlations within subsamples are constant. To illustrate how the observed relationship would behave in a perfectly constant environment, we simulate daily DCC correlations under the restriction that the true underlying errors $(e_{1,t}, e_{2,t})'$ remain unchanged over time. This situation is shown in the left graph of Panel B. By construction, the theoretical curve now perfectly fits the point cloud. We note some interesting similarities between the humped shaped relationship derived under the controlled simula-

tion with the one observed in the actual data. First, the general form, with high fluctuations in ρ when assets are uncorrelated and decreasing volatility as ρ approaches ± 1 , are clearly visible in the actual financial data. Second, the uncertainty concerning $\hat{\sigma}_{\hat{\rho}}$ is highest for uncorrelated assets. We interpret these similarities as an additional indicator that correlations in financial data are

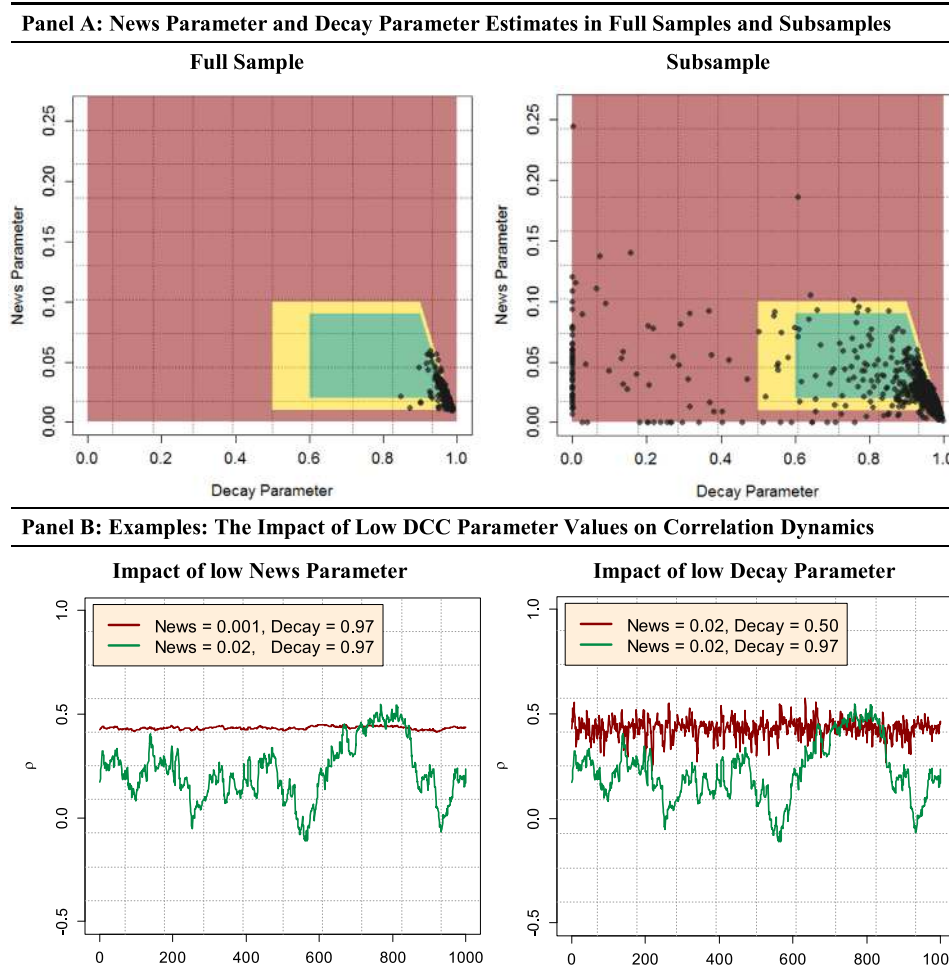


Fig. 8. Estimates of news parameter and decay parameter are different in subsamples. This figure shows the distribution of the news and decay parameter in full samples and subsamples. When parameter estimates move outside a narrow range they generate correlations that show little variation over time. For instance, a low news parameter generates near constant correlations. A low decay parameter generates correlations that fluctuate at low volatility around a constant value.

likely to be constant when structural breaks in correlations are accounted for.²¹

Our findings concerning the volatility in $\hat{\rho}_t$ do not extend to univariate GARCH models. We can repeat our analysis based on univariate GARCH volatility $\hat{\sigma}_t$ instead of $\hat{\rho}_t$. To collect observations on the level and volatility of $\hat{\sigma}$, we apply a volatility breakpoint test and measure the average sample volatility $\hat{\sigma}$ and the volatility of the GARCH volatility $\hat{\sigma}_{\hat{\sigma}}$ within subsamples. For detecting volatility breaks, we apply the test developed in [Inclán and Tiao \(1994\)](#).²² The right graph in Panel B shows that the relationship for $\hat{\sigma}$ appears linear and is strictly positive: higher levels of σ are associated with higher fluctuations in $\hat{\sigma}$. Hence, our findings concerning the behavior of correlations differ from previous results for univariate volatility models (e.g., [Rapach and Strauss, 2008](#)).

The humped shape function described by $\kappa(\rho)$ is a general result. However, it cannot show how breaks in the correlation structure transmit to DCC parameter values and ultimately affect the

volatility of $\hat{\rho}$. The link between correlation breaks and $\hat{\sigma}_{\hat{\rho}}$ is established through the model multiplier γ : correlation breaks distort correlation parameter estimates which in turn directly determine γ . In [Section 3](#) we found only small average changes in estimates of the news parameter and the decay parameter. Comparing full samples and subsamples, the average decay parameter decreased from 0.970 to 0.836, and the average news parameter actually increased slightly from 0.021 to 0.027. At first glance, none of these changes appear to be particularly large. However, we will show in the following that even small parameter changes can have a substantial impact on $\hat{\sigma}_{\hat{\rho}}$ if they move outside a narrow range. In the extreme case when the news and decay parameter estimates are close to zero, the generated correlations $\hat{\rho}_t$ are constant. We find this to be the case in a number of subsamples. To compare DCC parameter estimates in full samples from those obtained in subsamples we proceed in two steps. We first concentrate on the full sample and select asset correlations from which typical estimates of news and decay parameter values can be obtained. As before, the full samples contain daily data from 01/03/2000 to 12/31/2014 (3914 obs.). We then take the same assets but re-estimate the DCC parameters over the subsamples, where we follow our previous approach and define subsamples to lie on both sides of a correlation break point. Panel A in [Fig. 8](#) shows pairs of news and decay parameter estimates. The full sample estimates in the left graph were selected to be within a range of $0.01 < a < 0.06$ and

²¹ Of course, some differences are to be expected given that in practice, underlying correlations are not perfectly constant and are subject to differences DCC parameters and sample size.

²² [Rapach and Strauss \(2008\)](#) propose an adjusted Inclán-Tiao test that is more effective in the presence of GARCH effects. For our purpose, the adjustment is of minor importance since our primary goal is to collect data on the level and volatility of $\hat{\sigma}$, and not to detect the correct breakpoint date.

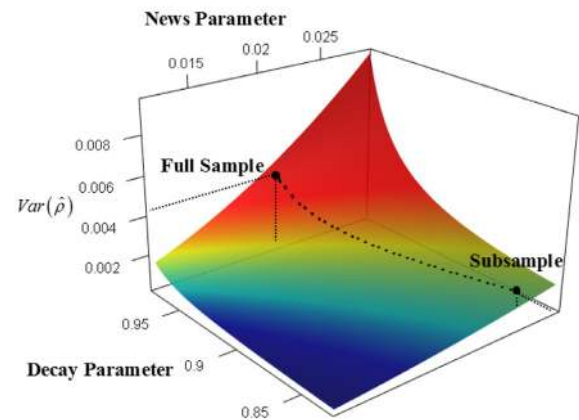
$0.8 < b < 0.99$. Only parameter combinations that lie within this range generate typical dynamic correlation. The range is indicated by a green background to highlight that the area of allowed parameters is small relative to the theoretically possible parameter space. The right graph shows the distribution of DCC estimates when re-estimated over subsamples. The estimates are based on the same asset pairs that were used in the full sample analysis and are therefore directly comparable. We exclude subsamples with less than 500 observations to remove small sample effects on the parameters. A large fraction of the parameter estimates now lie outside the range of typical parameter values. A number of decay parameter values are smaller than 0.5 and some are even zero.

In Panel B of Fig. 8 we illustrate how parameter estimates outside the normal range impact the correlation dynamics. The left graph shows a generated correlation when the news parameter is close to zero. As expected, the generated correlations are constant. The right graph shows the situation when the news parameter takes on a typical value but the decay parameter is 0.5. The generated correlation is quasi constant with small fluctuations around a straight line. We can summarize our findings concerning parameter changes as follows: The typical dynamic pattern of $\hat{\rho}_t$ that has become a stylized fact for many financial assets requires that both, the news and the decay parameter lie within a narrow range.²³ Parameter values outside this range lead to correlations that lack typical dynamics. When estimated over subsamples that do not contain a break, many of the correlations turned out to be constant. Our findings suggest that many correlation dynamics are spurious and disappear once correlation breaks are controlled for.

It is instructive to investigate how the dynamics in $\hat{\rho}_t$ respond to the interaction of the news and the decay parameter. We show that this interaction is strongly nonlinear so that seemingly small changes in model parameters can have significant effects. In particular, the observed decrease in the decay parameter from 0.970 in full samples to 0.836 in subsamples has important implications for the dynamics of $\hat{\rho}_t$. Panel A of Fig. 9 shows the location of full sample and subsample DCC parameters on the $Var(\hat{\rho}_t^*)$ surface. The surface shows a distinct upward slope for high news and decay parameter values. The full sample estimates with news parameter $a = 0.021$, decay parameter $b = 0.970$, and hence $a + b = 0.991$ constitute one such combination that generate large fluctuations in $Var(\hat{\rho}_t^*)$. For the full sample, we find $Var(\hat{\rho}_t^*)$ to be 0.004. The upward slope illustrates the “narrow band” that we mentioned before. On the other hand, $Var(\hat{\rho}_t^*)$ quickly converges towards zero if one DCC parameter decreases just slightly. The location of the subsample DCC parameters shows a marginally higher news parameter value of 0.027 that is more than compensated by a significantly lower decay parameter value of 0.836. The sum of both DCC parameters is 0.863 and therefore significantly below one. The location on the surface indicates that these values produce much less volatile estimate of $\hat{\rho}_t$. In fact, our estimate for $Var(\hat{\rho}_t^*)$ in subsamples is just 0.001, a quarter of its full sample size. This supports the notion that DCC correlations are constant in subsamples. From the findings in Panel A, we conclude that due to the nonlinear interaction of DCC parameters, even small deviation can have important implications for the variation in $\hat{\rho}_t$.

Panel B of Fig. 9 shows how the findings concerning the variance of $\hat{\rho}_t$ also extend to the autocovariance and therefore the dynamics of $\hat{\rho}_t$. To highlight the impact of the decay parameter on various autocovariance lags s , the value for the news parameter is fixed at 0.02. The shape of the autocovariance surface shows that $\hat{\rho}_t$ has very short or no memory for most parameter combinations. Significant dynamics only emerge for decay parameter values that are large, so that the sum of a and b are close to one. Again, the

Panel A: Subsample DCC Estimates Imply Lower Fluctuations in $\hat{\rho}$



Panel B: Subsample Decay Parameter Values Imply Lower Persistence in $\hat{\rho}$

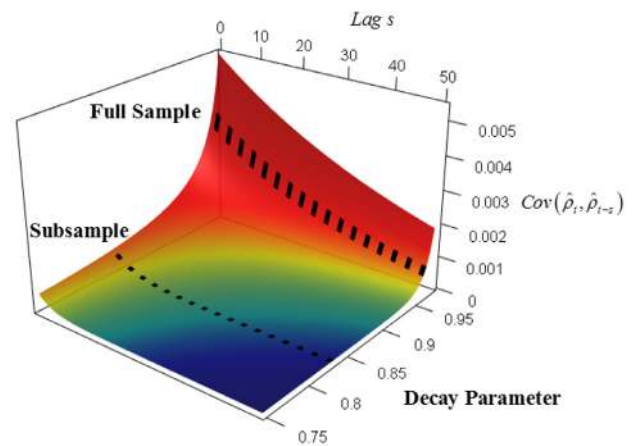


Fig. 9. The nonlinear impact of DCC parameters on $Var(\hat{\rho})$. This figure shows that the news and decay parameter values need to lie within a narrow area in order to produce meaningful dynamics in $\hat{\rho}$. The parameter estimates found in subsamples are not sufficiently large to have either significant fluctuations or noticeable autocovariance.

average parameter location in full samples is sufficiently close to one to produce the distinct dynamics that are typical for correlation estimates of many financial time series. In contrast, average subsample DCC parameters produce correlations with little serial correlation. The findings in Panels A and B indicate that when controlling for breaks, dynamic conditional correlation estimates show little variations and no significant dynamics. The implication is that the nature of correlations is constant.

5. Volatility ratios

Like variances, correlations are unobserved and need to be estimated. The absence of a true observed correlation complicates model comparison in empirical applications. Engle and Colacito (2006) propose a method that allows for an effective comparison of correlation models within the portfolio setting. The idea is to use portfolio variance as a measure of the effectiveness of a dynamic correlation model where lower variance indicates a better correlation model. Consider the standard portfolio optimization problem

$$\min_{w_{i,t}} w'_{i,t} H_{i,t} w_{i,t} \quad s.t. \quad w'_{i,t} \mu = \mu_0, \tag{13}$$

²³ The results also hold in the case of the EWMA model for various λ .

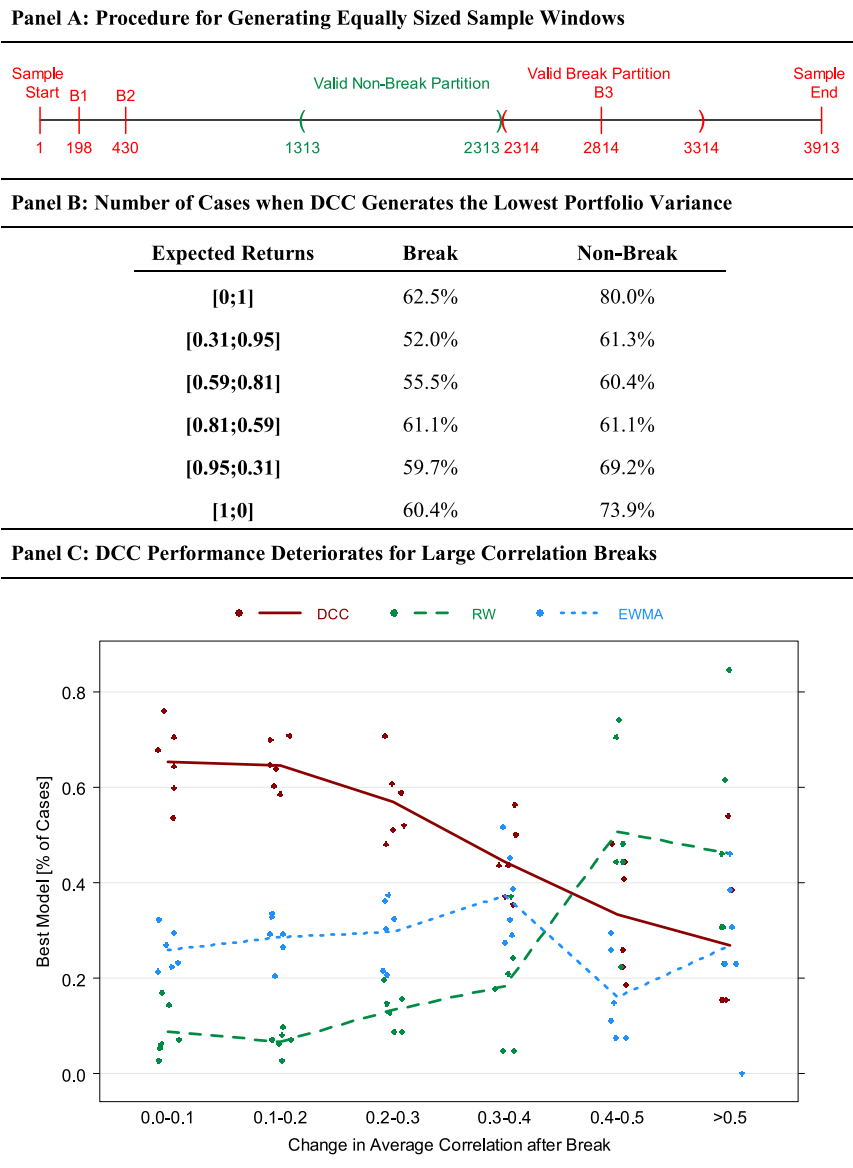


Fig. 10. DCC model performance in asset portfolios.

This figure shows the relative performance of asset portfolios based on different dynamic covariance matrix estimators. Panel A shows how break and non-break samples with equal size are selected. Panel B shows that the DCC model performs less well in samples containing a correlation break. Panel C shows that when correlation breaks are large, a simple rolling window estimator is the best choice.

where $\mathbf{H}_{i,t}$ is the conditional covariance matrix of model i at time t , μ is the vector of portfolio weights and $\mu_0 > 0$ is the required target return. Campbell et al. (1997) show that the solution and therefore the optimal portfolio weights can be estimated by $\mathbf{w}_{i,t} = (\mathbf{H}_{i,t}^{-1}\mu)/(\mu'\mathbf{H}_{i,t}^{-1}\mu)\mu_0$. The volatility of the portfolio return $\mathbf{w}'_{i,t}\mathbf{r}_t$ can then be obtained as $\hat{\sigma}_{i,t} = \sqrt{\mathbf{w}'_{i,t}\mathbf{H}_{i,t}\mathbf{w}_{i,t}}$. An efficiently estimated dynamic covariance matrix $\mathbf{H}_{i,t}$ will be reflected in a low portfolio volatility $\hat{\sigma}_{i,t}$. A comparison of covariance matrix estimators based on this approach has been applied for instance in DeMiguel et al. (2009) but was entirely based on constant moments.

To compare portfolio volatilities Engle and Colacito (2006) form volatility ratios $VR_{i,t}$. In the ideal case, where the true covariance matrix $\mathbf{\Omega}_t$ is known, the volatility ratio would be

$$VR_{i,t} = \frac{\sqrt{\mathbf{w}'_{i,t,t}\mathbf{H}_{i,t}\mathbf{w}_{i,t}}}{\sqrt{\mathbf{w}'_t\mathbf{\Omega}_t\mathbf{w}_t}} \quad (14)$$

By construction, the volatility ratio in Eq. (14) is larger than one and shows an excess portfolio volatility that is based on the estimate $\mathbf{H}_{i,t}$ rather than the true correlation matrix $\mathbf{\Omega}_t$. The extent to which $VR_{i,t}$ exceeds one indicates the inefficiency of the dynamic covariance estimator. Since true correlations are unobservable in practice, Engle and Colacito (2006) consider a version of $VR_{i,t}$ that is based on the different model specifications for $\mathbf{H}_{i,t}$:

$$\widehat{VR}_{i,t} = \frac{\hat{\sigma}_{i,t}}{\min(\hat{\sigma}_{1,t}, \dots, \hat{\sigma}_{n,t})} \quad (15)$$

We follow this approach and consider the $n = 3$ correlation models DCC, RW, and EWMA. To evaluate the performance of our three models in terms of portfolio variance we use the same 40 asset data set that was employed in previous sections. Our focus is on the model performance in samples that contain breaks and more stable situations that do not contain any breaks. One potential issue is the different sample size that can influence estimation

results. For instance, a larger non-break sample could appear to produce smaller portfolio volatilities because the large number of observations leads to more precise model parameters. To circumvent this problem we select break and non-break samples of equal size. To illustrate this issue further, Panel A of Fig. 10 shows the sample selection process for the correlation between the stock indices of Italy (MIB) and the Netherlands (AEX). For this sample, we detect a correlation breakpoint at observation 198, another at observation 430, and a third at observation 2814. The third observation is a suitable break partition and we select a 1000 observation window from observation 2314 to 3314 with the breakpoint in the center of that window. The data prior to the break window contain sufficient observations to select another 1000 non-break observation window for comparison. We analyze all other asset pairs in this way and collect 429 break and non-break partitions of size 1000.

The portfolio variance is not only a function of the estimated covariance matrix but also depends on the expected return vector μ that enters the optimization as an input. Although the variance does not appear to respond very strongly to this assumption, we follow Engle and Colacito (2006) and compute volatility ratios over a range of expected returns. Panel B of Fig. 10 shows the percentage of cases in which the smallest portfolio variance is generated by the DCC model and hence has a volatility ratio of one. Among the three correlation models, the DCC model performs well in situations that do not contain a structural break producing the lowest portfolio variance in 60% to 80% of all cases. In samples that contain a correlation break, the DCC parameters are biased and the portfolio performance decreases. Still, the overall performance of the DCC is quite remarkable compared to its competitors. However, the break samples that form the basis for the results in Panel B only test for the presence of a statistically significant break whereas the DCC parameters are likely to respond also to the size of the break. In Panel C of Fig. 10, we look at model performance conditioning on the size of the break indicated on the x-axis of the graph. For economically small correlation breaks up to 0.4 the DCC model produces lower portfolio variances than the rolling window or the EWMA model. However, the performance of the DCC model deteriorates quickly as the break size increases. For instance, for small breaks of less than 0.1, DCC is the best model in 65% of all cases. For larger breaks between 0.4 and 0.5, the DCC is the best choice in 33% of all cases. For large correlation breaks of more than 0.5, the distorting effect on the DCC correlation is so strong that the simple rolling window estimator leads to better portfolio performance. Multivariate GARCH models like DCC are often praised for their dynamic flexibility to accommodate changes in the return pattern. Our results indicate that in the presence of correlation breaks a rolling window estimator may perform better despite its simplicity.

6. Conclusion

In this paper, we provide empirical evidence that daily correlation dynamics among financial assets are spurious. The typical correlation dynamics that can be observed in the data are a direct consequence of correlation breaks that occur in response to financial and economic shocks. The presence of breaks affects the correlation news parameter a and the decay parameter b . The news parameter, which measures the response to shocks, and the decay parameter, which measures the persistence of correlations, interact in a nonlinear way to generate the correlation dynamics that we usually observe for financial assets. Once these breaks are controlled for, the parameters driving the correlation dynamics change in important ways. A number of parameters are now close to zero and generate constant correlations. The average estimate of the decay parameter which is upward biased in the presence of correla-

tion breaks decreases in subsamples. The sum of a and b which is usually found to be close to one is therefore lower. These subtle changes remove the parameters from an area of influence that generate typical correlation dynamics. The variance and autocovariance estimates are now lower, indicating that correlation estimates fluctuate at low volatility around a straight line. The true nature of correlations is therefore likely to be constant. The implication for empirical correlation estimates is that the path generated by multivariate GARCH correlations should be interpreted with caution. A portfolio spanning the main asset classes is shown to respond to the way correlations are estimated. A significant break in the correlation structure can distort DCC correlations and lead to a higher portfolio variance. We show that investors can resort to simple solutions such as a rolling window estimator when updating their portfolio weights.

In summary, our results provide a rationale for the often controversial discussion of the value added of dynamic conditional correlation models. A number of dynamic correlation models have formed the basis for a significant amount of important research and we do not propose to reject these models entirely. However, academics and practitioner should be aware of the practical limitations of these models that arise in many finance applications.

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at [doi:10.1016/j.jbankfin.2017.07.003](https://doi.org/10.1016/j.jbankfin.2017.07.003).

References

- Adams, Z., Glück, T., 2015. Financialization in commodity markets: a passing trend or the new normal? *J. Bank. Finance* 60, 93–111.
- Adrian, T., Brunnermeier, M., 2016. CoVaR. *Am. Econ. Rev.* 106 (7), 1705–1741.
- Aielli, G., 2013. Dynamic conditional correlation: On properties and estimation. *J. Bus. Econ. Stat.* 31 (3), 282–299.
- Andreou, E., Ghysels, E., 2002. Detecting multiple breaks in financial market volatility dynamics. *J. Appl. Econometrics* 17 (5), 579–600.
- Andreou, E., Ghysels, E., 2003. Tests for breaks in the conditional co-movement of asset returns. *StatisticaSinica* 13 (4), 1045–1073.
- Aslanidis, N., Martinez, O., 2016. A Threshold Conditional Correlation Approach to Modeling Equity and Bond Returns. University of Rovira Virgili Working Paper.
- Bauwens, L., Laurent, S., Rombouts, J., 2006. Multivariate GARCH models: a survey. *J. Appl. Econometrics* 21 (1), 79–109.
- Bera, A., Kim, S., 2002. Testing constancy of correlation and other specifications of the BGARCH model with an application to international equity returns. *J. Empirical Finance* 9 (2), 171–195.
- Bollerslev, T., 1990. Modelling the coherence in short-run nominal exchange rates: a multivariate generalized ARCH model. *Rev. Econ. Stat.* 72 (3), 498–505.
- Bollerslev, T., Engle, R., Wooldridge, 1988. A capital asset pricing model with time-varying covariances. *J. Polit. Econ.* 96 (1), 116–131.
- Bollerslev, T., Chou, R., Kroner, K., 1992. ARCH modeling in finance: a review of theory and empirical evidence. *J. Econometrics* 52 (1–2), 5–59.
- Brenner, M., Pasquariello, R., Subrahmanyam, M., 2009. On the volatility and co-movement of U.S. financial markets around macroeconomic news announcements. *J. Financ. Quant. Anal.* 44 (6), 1265–1289.
- Cai, Y., Chou, R.Y., Li, D., 2009. Explaining international stock correlations with CPI fluctuations and market volatility. *J. Bank. Finance* 33 (11), 2026–2035.
- Campbell, J.W., Lo, A.W., MacKinlay, C., 1997. *The Econometrics of Financial Markets*. Princeton University Press.
- Caporin, M., McAleer, M., 2008. Scalar BEKK and indirect DCC. *J. Forecast.* 27 (6), 537–549.
- Caporin, M., McAleer, M., 2012. Do we really need both BEKK and DCC? A tale of two multivariate GARCH models. *J. Econ. Surv.* 26 (4), 736–751.
- Cappiello, L., Engle, R.F., Sheppard, K., 2006. Asymmetric dynamics in the correlations of global equity and bond returns. *J. Financ. Econometrics* 4 (4), 537–572.
- Chang, C., McAleer, M., Tansuchat, R., 2011. Crude oil hedging strategies using dynamic multivariate GARCH. *Energy Econ.* 22 (5), 912–923.
- Cheng, I.-H., Xiong, W., 2014. Financialization of commodity markets. *Annu. Rev. Financ. Econ.* 6, 419–441.
- Christiansen, C., 2000. Macroeconomic announcement effects on the covariance structure of government bond returns. *J. Emp. Finance* 7 (5), 479–507.
- Davidson, J., 1994. *Stochastic Limit Theory: An Introduction for Econometricians*. Oxford University Press.
- DeMiguel, V., Garlappi, L., Uppal, R., 2009. Optimal versus naïve diversification: how inefficient is the 1/N portfolio strategy? *Rev. Financ. Stud.* 22 (5), 1915–1953.
- Embrechts, P., McNeil, A., Straumann, D., 1999. Correlation: pitfalls and alternatives. *Risk* 12, 69–71.

- Engle, R., 2002. Dynamic conditional correlation – a simple class of multivariate GARCH models. *J. Bus. Econ. Stat.* 20 (3), 339–350.
- Engle, R., Colacito, R., 2006. Testing and valuing dynamic correlations for asset allocation. *J. Bus. Econ. Stat.* 24 (2), 238–253.
- Engle, R., Granger, C., Kraft, D., 1984. Combining competing forecasts of inflation based on a multivariate ARCH model. *J. Econ. Dyn. Control* 8 (2), 151–165.
- Engle, R., Kroner, K., 1995. Multivariate simultaneous generalized ARCH. *Econometric Theory* 11 (1), 122–150.
- Engle, R., Ng, V., 1993. Measuring and testing the impact of news on volatility. *J. Finance* 48 (5), 1749–1778.
- Engle, R., Sheppard, K., 2001. Theoretical and Empirical Properties of Dynamic Conditional Correlation Multivariate GARCH. NBER Working Papers 8554. National Bureau of Economic Research.
- Forbes, K.J., Rigobon, R., 2002. No contagion, only interdependence: measuring stock market comovements. *J. Finance* 57, 2223–2261.
- Foster, D., Nelson, D., 1996. Continuous record asymptotics for rolling sample variance estimators. *Econometrica* 64 (1), 139–174.
- Fleming, K., Kirby, C., Ostdiek, B., 2001. The economic value of volatility timing. *J. Finance* 56 (1), 329–352.
- Galeano, P., Wied, D., 2014. Multiple break detection in the correlation structure of random variables. *Comput. Stat. Data Anal.* 76, 262–282.
- Hale, G., Obstfeld, M., 2016. The euro and the geography of international debt flows. *J. Eur. Econ. Assoc.* 14 (1), 115–144.
- Hamilton, J., Susmel, R., 1994. Autoregressive heteroscedasticity and changes in regime. *J. Econometrics* 64 (1), 307–333.
- Heaney, R., Srikanthakumar, S., 2012. Time-varying correlation between stock market returns and real estate returns. *J. Emp. Finance* 19 (4), 583–594.
- Henderson, B.J., Pearson, N.D., Wang, L., 2015. New evidence on the financialization of commodity markets. *Rev. Financ. Stud.* 28 (5), 1285–1311.
- Hillebrand, E., 2005. Neglecting parameter changes in GARCH models. *J. Econometrics* 129 (1), 121–138.
- Inclán, C., Tiao, G.C., 1994. Use of Cumulative sums of squares for retrospective detection of changes in variance. *J. Am. Stat. Assoc.* 89, 913–923.
- Karolyi, A., Stulz, R., 1996. Why do markets move together? An investigation of U.S.-Japan stock return comovements. *J. Finance* 51 (3), 951–986.
- King, M., Sentana, E., Wadhvani, S., 1994. Volatility and links between national stock markets. *Econometrica* 62 (4), 901–933.
- Kwan, C., 2008. Estimation error in the average correlation of security returns and shrinkage estimation of covariance and correlation matrices. *Finance Res. Lett.* 5 (4), 236–244.
- Lane, P.R., 2012. The European sovereign debt crisis. *J. Econ. Perspect.* 26 (3), 49–68.
- Longerstaey, J., More, L., 1995. Introduction to RiskMetrics™, fourth ed. Morgan Guaranty Trust Company, New York.
- Longin, F., Solnik, B., 1995. Is the correlation in international equity returns constant: 1960–1990? *J. Int. Money Finance* 14 (1), 3–26.
- Lütkepohl, H., 2006. New Introduction to Multiple Time Series Analysis. Springer, Berlin Heidelberg.
- Mittnik, S., Paolella, M.S., 2000. Conditional density and value-at-risk prediction of Asian currency exchange rates. *J. Forecast.* 19 (4), 313–333.
- Mody, A., Sandri, D., 2012. The Eurozone crisis: how banks and sovereigns came to be joined at the hip. *Econ. Policy* 27 (70), 199–230.
- Moskowitz, T.J., 2003. An analysis of covariance risk and pricing anomalies. *Rev. Financ. Stud.* 16 (2), 417–457.
- Ofek, E., Richardson, M., 2003. DotCom mania: the rise and fall of internet stock prices. *J. Finance* 58 (3), 1113–1138.
- Pakel, C., Shephard, N., Sheppard, K., Engle, R.F., 2014. Fitting Vast Dimensional Time-Varying Covariance Models Working paper available at <https://www.aeaweb.org/conference/2015/retrieve.php?pdfid=101>.
- Pelletier, D., 2006. Regime-switching for dynamic correlations. *J. Econometrics* 131 (1–2), 445–473.
- Pérignon, C., Smith, D., 2010. Diversification and value-at-risk. *J. Bank. Finance* 34 (1), 55–66.
- Pukthuanthong, K., Roll, R., 2011. Gold and the dollar (and the Euro, Pound, and Yen). *J. Bank. Finance* 35 (8), 2070–2083.
- Rapach, D., Strauss, J., 2008. Structural breaks and GARCH models exchange rate volatility. *J. Appl. Econometrics* 23 (1), 65–90.
- Sadorsky, P., 2012. Correlations and volatility spillovers between oil prices and the stock prices of clean energy and technology companies. *Energy Econ.* 34 (1), 248–255.
- Silvennoinen, A., Teräsvirta, T., 2009. Modeling multivariate autoregressive conditional heteroscedasticity with the double smooth transition conditional correlation GARCH Model. *J. Financ. Econ.* 7 (4), 373–411.
- Tang, G.Y.N., 1995. Intertemporal stability in international stock market relationships: a revisit. *Q. Rev. Econ. Finance* 35 (1), 579–593.
- Tse, Y., Tsui, A., 1999. A note on diagnosing multivariate conditional heteroscedasticity models. *J. Time Ser. Anal.* 20 (6), 679–691.
- Tse, Y., 2000. A test for constant correlations in a multivariate GARCH model. *J. Econometrics* 98 (2), 107–127.
- West, K.D., Cho, D., 1995. The predictive ability of several models of exchange rate volatility. *J. Econometrics* 69 (2), 367–391.
- Wied, D., Krämer, W., Dehling, H., 2012. Testing for a change in correlation at an unknown point in time using an extended functional delta method. *Econometric Theory* 28 (3), 1–20.
- You, L., Daigler, R., 2010. Is international diversification really beneficial? *J. Bank. Finance* 34 (1), 163–173.