Journal of Banking & Finance 48 (2014) 1–12

Contents lists available at ScienceDirect

Journal of Banking & Finance

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

The role of correlation dynamics in sector allocation ☆

Elena Kalotychou [a,b,*], Sotiris K. Staikouras [a,b,1], Gang Zhao [c,2]

[a] Cass Business School, City University, London EC1Y 8TZ, UK
[b] ALBA Graduate Business School, Athens 16671, Greece
[c] School of Management, University of Bath, Bath BA2 7AY, UK

ARTICLE INFO

Article history:
Received 1 July 2013
Accepted 29 June 2014
Available online 21 July 2014

JEL classification:
C32
C52
C53
F21
G11
G15

Keywords:
Correlation timing
Asymmetry
Structural break
Portfolio performance
Transaction costs

ABSTRACT

This paper assesses the economic value of modeling conditional correlations for mean–variance portfolio optimization. Using sector returns in three major markets we show that the predictability of models describing empirical regularities in correlations such as time-variation, asymmetry and structural breaks leads to significant performance gains over the static covariance strategy. Investors would be willing to pay a fee of up to 983 basis points to switch from the static to the dynamic correlation portfolio and about 100 basis points more for capturing asymmetries and shifts in correlations. The gains are robust to the crisis, transaction costs and are most pronounced for monthly rebalancing.

1. Introduction

Volatility and correlation among asset returns are central to portfolio allocation and risk management. A burgeoning literature in financial economics has focused on time series models for asset return volatility and their comovement. Various Multivariate Generalized Autoregressive Conditional Heteroskedasitcity (MGARCH) models, such as the Dynamic Conditional Correlation (DCC) model of Engle (2002), have been developed to capture the well-documented time variation in correlations and other dynamic aspects of comovement between financial risks.

Correlation asymmetry is one regularity that has been widely found in the second moment of equity returns although the economic rationale behind the clustering of bad news is relatively less researched. Longin and Solnik (2001) show that correlations rise in bear markets. Ang and Bekaert (2002) document the presence of a high volatility-high correlation regime in the US, UK and Germany, which coincides with a bear market and refutes the benefits of international diversification. Cappiello et al. (2006) find support for asymmetry in the correlations of international equity and bond returns, while Bekaert et al. (2005) attribute jumps in cross-market correlations during crises to dependence on a common factor.

Structural breaks have also been documented in correlations and can have a fundamental impact on global markets. Billio and Pelizzon (2003) find that correlations of European markets increased following the European Monetary Union (EMU). Longin and Solnik (2001) suggest that the level and structure of global correlations shifted considerably over time. Cappiello et al. (2006) find significant correlation rise post-EMU not mirrored in conditional volatility indicating greater market integration.

☆ The authors would like to thank participants at the IFABS Conference 2012, Valencia, the FMA European Meeting 2010, Hamburg, the FMA Annual Meeting 2009, Reno, the Multinational Finance Society Conference 2009, Crete, the FMA European Meeting 2009, Turin, the British Accounting Association Annual Conference 2009, Dundee, the 28th International Symposium on Forecasting 2008, Nice, and seminar participants at Cass Business School for useful comments on earlier versions of the paper. Special thanks to Spyros Mesomeris, Director of the Deutsche Bank Quantitative Strategy Team for constructive discussions and to the two anonymous referees for their helpful comments. The usual disclaimer applies.

* Corresponding author at: Cass Business School, City University, London EC1Y 8TZ, UK. Tel.: +44 20 7040 5259.

E-mail addresses: e.kalotychou@city.ac.uk (E. Kalotychou), sks@city.ac.uk (S.K. Staikouras), gz249@bath.ac.uk (G. Zhao).

[1] Tel.: +44 20 7040 0165.
[2] Tel.: +44 1225 383968.

http://dx.doi.org/10.1016/j.jbankfin.2014.06.025
0378-4266/

There has been growing consensus that employing static long-term historical relationships between assets in portfolio management may lead to substantial underperformance in the face of increased market volatility, changing correlations and frequent regime shifts. This study assesses the economic merit of forecasting return correlation dynamics for sector allocation. We seek to generate profitable trading strategies through correlation predictability, that is, correlation timing, a notion introduced by Engle and Colacito (2006). The contribution to the literature is twofold. First, we investigate the economic value of capturing stylized facts of asset correlations such as time variation, asymmetry and structural breaks. To do so we employ a dynamic mean–variance framework, which incorporates investor risk aversion, transaction costs and different rebalancing frequencies. Second, as the value and viability of market timing strategies during the recent financial crisis has often been questioned, we empirically examine the benefits of correlation timing over the crisis period (2007–2009) and its aftermath (2009–2012).

The pertinent empirical literature mainly focuses on the economic value of volatility timing (Fleming et al., 2001, 2003; Della Corte et al., 2009). The evaluation of conditional correlation estimators has largely focused on statistical metrics and less attention has been paid to the economic value of capturing the empirical regularities in correlations. Engle and Sheppard (2001) show that the DCC model outperforms the industry standard RiskMetrics exponential smoother on the basis of residual normality and lower portfolio standard deviations. Engle and Colacito (2006) show that the efficiency loss of mean–variance portfolios decreases with correlation accuracy and that assuming constant correlation during volatile correlation phases is costly. But important issues such as the profitability of correlation predictability, the impact of transaction costs on active allocation and the value of the latter during market downturns or for different risk-aversions have not been examined as yet.[3]

Our analysis is based on daily prices from ten sector indices in three major markets (Japan, UK, US) over July 1996–April 2012. The findings suggest that correlation timing is fruitful to sector investors. Dynamic correlation strategies deliver significant out-of-sample gains in risk-adjusted returns, which are more pronounced for monthly rebalancing and are robust to reasonable transaction costs. Risk-averse investors are willing to pay a fee of up to 983 basis points (bp) to switch from the static covariance portfolio to the dynamic DCC portfolio and up to an additional 100 bp to also account for correlation asymmetries and regime shifts. The Sharpe Ratio (*SR*) accrual of dynamic portfolios can be as high as 0.48 and rises a further 0.08 when asymmetries and structural breaks are captured. Exploiting correlation dynamics appears more beneficial during the crisis: risk-adjusted returns rise to 0.60 in excess of the static portfolio and performance fees largely increase.

The remainder of the paper is organized as follows. Section 2 describes the data. Section 3 delineates the conditional correlation models and the performance evaluation framework. Section 4 presents the empirical results. Section 5 concludes.

2. Data

The empirical analysis is based on daily prices for ten sector indices from the Nikkei 225, FTSE-All and S&P500 obtained from Thomson Reuters DataStream International, namely, Energy (ENG), Basic Material (BML), Industrial (IND), Consumer Goods (CGS), Health Care (HCR), Consumer Service (CSV), Telecommunication (TEL), Utility (UTL), Financial (FIN) and Technology (TEC). The sample spans the period from July 1, 1996 to April 30, 2012, which amounts to a total of around 3900 daily logarithmic returns (in local currency) for each sector portfolio. The three-month Japanese interbank loan rate, the UK LIBOR, and the US Treasury bill rate proxy the risk free asset. The descriptive statistics in Table 1 show positive mean daily returns for most sectors.

All daily returns are non-normally distributed, particularly in the form of leptokurtosis. The extent and direction of skewness differs across sectors and equity markets. Most of the sector returns in the three markets are significantly negatively skewed. The Augmented Dickey–Fuller (ADF) test strongly rejects the hypothesis of a unit root for all return series. The Ljung–Box *Q*-statistic on raw/squared daily returns portrays serial dependence in all sectors. The strong evidence of volatility clustering supports the stylized fact that there is far more predictability in conditional volatility than in returns.

The analysis is based on domestic sector portfolios in each of the three markets and so within-country sector correlations are of relevance. The unconditional sector correlations over the sample period are significantly positive. The average sector correlation within Japan, UK and US is 59.7%, 48.1% and 62.9%, respectively.[4] Consumer services and industrials exhibit the highest correlation with other sectors in their respective markets at 65.8% and 64.7%, while utilities are the least correlated.

Our empirical framework is designed to assess the economic differences materializing from rival correlation forecasting approaches. The sample is divided into an in-sample estimation period from July 1, 1996 to June 30, 2005 ($T = 2274$, 2266, 2209 days, respectively, for the Japanese, UK and US sector portfolios) and a holdout evaluation period from July 1, 2005 to April 30, 2012 ($T^* = 1676$, 1727, 1720 days, respectively, for the three domestic sector portfolios). The choice of out-of-sample period enables us to evaluate the performance of correlation timing over three distinctive phases of the recent global financial crisis, i.e. the pre-crisis (July 2005–July 2007), crisis (August 2007–February 2009), and post-crisis (March 2009–April 2012) periods. The conditional correlation models are re-estimated over a rolling window of length-T to generate one-step-ahead covariance matrix forecasts.[5]

3. Methodology

The analysis builds upon the recursive construction of optimal mean–variance sector portfolios in the Japanese, UK and US markets and their out-of-sample performance based on incremental utility and risk-adjusted returns. For this purpose daily sector correlation and volatility forecasts, the main inputs alongside expected returns for active mean–variance allocation, are generated using the models outlined below.

3.1. The conditional covariance structure

Let r_t denote the day t logarithmic close-to-close return vector on n risky assets and ξ_{t-1} be the information set available at the end of day t−1. The $[n \times 1]$ conditional expected return vector of r_t is defined as $\mu_t \equiv \mu_{t|t-1} = \mathrm{E}[r_t|\xi_{t-1}]$, while $H_t \equiv H_{t|t-1} = \mathrm{E}[(r_t - \mu_t)$

[3] While DeMiguel et al. (2009) argue that the naïve *1/N* diversification strategy is able to outperform the mean–variance asset allocation, their findings have been questioned by Kirby and Ostdiek (2012) who document that active mean–variance timing is superior to naïve diversification but can be severely affected by transaction costs.

[4] The three mean correlations are strongly significant with *t*-statistics 58.7, 39.6 and 65.8. The *t*-statistic is computed as $\rho\sqrt{(T-2)/(1-\rho^2)}$ and follows a Student-*t* distribution with $(T-2)$ degrees of freedom.

[5] According to Clark and McCracken (2001), the ratio between the out-of-sample and in-sample period observations (π) should not be too large or small. In the current study, π is ranging from 0.74 to 0.78, thereby leaving a sizeable number of observations in each of the in-sample and out-of-sample portions.