Economic Modelling 45 (2015) 50–73

Contents lists available at ScienceDirect

Economic Modelling

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

Modelling stock return volatility dynamics in selected African markets

Daniel King, Ferdi Botha *

Department of Economics and Economic History, Rhodes University, Grahamstown, South Africa

ARTICLE INFO

Article history:
Accepted 6 November 2014
Available online xxxx

Keywords:
Stock returns
Volatility
GARCH
Africa

ABSTRACT

This paper examines whether accounting for structural changes in the conditional variance process, through the use of Markov-switching models, improves estimates and forecasts of stock return volatility over those of the more conventional single-state (G)ARCH models, within and across selected African markets for the period 2002–2012. In the univariate portion of the paper, the performances of various Markov-switching models are tested against a single-state benchmark model through the use of in-sample goodness-of-fit and predictive ability measures. In the multivariate context, the single-state and Markov-switching models are comparatively assessed according to their usefulness in constructing optimal stock portfolios. Accounting for structural breaks in the conditional variance process, conventional GARCH effects remain important in capturing heteroscedasticity. However, those univariate models including a GARCH term perform comparatively poorly when used for forecasting purposes. In the multivariate study, the use of Markov-switching variance–covariance estimates improves risk-adjusted portfolio returns relative to portfolios constructed using the more conventional single-state models. While there is evidence that some Markov-switching models can provide better forecasts and higher risk-adjusted returns than those models which include GARCH effects, the inability of the simpler Markov-switching models to fully capture heteroscedasticity in the data remains problematic.

1. Introduction

Regardless of the many likely causes of stock market volatility (Abel, 1988; Adam et al., 2008; Aggarwal et al., 1999; Diebold and Yilmaz, 2007; Olsen, 1998; Schwert, 1989; Shiller, 1987), a reliable statistical model of stock return volatility is important for the pricing of equity derivative securities and effective hedging of stock market risk (Hamilton and Susmel, 1994; Wang and Theobald, 2008). In addition, changes in the co-movement of stock returns across international markets during high- and low-volatility periods have major implications for diversification strategies (Li, 2009; Ramchand and Susmel, 1998).

A common feature of Generalised Autoregressive Conditional Heteroscedasticity (GARCH)-type models using daily financial data is the high level of persistence attributed to shocks. Many GARCH studies involving financial series have found that an approximate unit root process generates the estimated variance (Engle and Bollerslev, 1986; Susmel, 1999). However, it has been shown that in the presence of structural breaks, GARCH-type models can impose a spuriously high level of persistence of shocks on volatility forecasts (Diebold, 1986; Lamoureux and Lastrapes, 1990; Timmerman, 2000). This finding has led to the parallel development of the Markov-switching ARCH (SWARCH) model, which allows for endogenously identified structural shifts in the volatility generating process (Cai, 1994; Hamilton and Susmel, 1994).

* Corresponding author at: P.O. Box 94, Rhodes University, Grahamstown 6140, South Africa. Tel.: +27 466038673.
E-mail address: f.botha@ru.ac.za (F. Botha).

Having controlled for regime shifts, the persistence of shocks on volatility forecasts is generally reduced in a statistically significant way (Cai, 1994; Edwards and Susmel, 2003; Hamilton and Susmel, 1994; Marcucci, 2005; Ramchand and Susmel, 1998). Thus, while volatility clustering is still captured, the regime-switching models allow this clustering to be generated by regime changes in addition to within-regime persistence of shocks. That is, where single-regime GARCH models imply a purely time-varying variance, regime-switching models allow for volatility that is both time-varying and state-varying (Ramchand and Susmel, 1998). This specification can thus offer a more intuitively appealing interpretation of the volatility-clustering phenomenon than the single-regime GARCH models, as well as improve forecasts due to a higher likelihood of stationarity.

It was initially believed that the regime-switching models would in practice have to be restricted to low order SWARCH due to the recursive nature of GARCH models and the resulting intractability of maximum likelihood estimation for studies with large samples (Cai, 1994; Hamilton and Susmel, 1994). However, this estimation problem has been largely addressed by Gray (1996). The use of Gray's (1996) Markov-switching GARCH (MS–GARCH) procedure or it's extensions (Dueker, 1997; Haas et al., 2004; Klaassen, 2002) has allowed richer comparison of parameter estimates across models, as it nests the popular GARCH(1,1) as a special case (Gray, 1996).

Markov-switching models often provide a better in-sample fit of the data or more accurate forecasts than the conventional single-state GARCH extensions (Bollen et al., 2000; Cai, 1994; Canarella and Pollard, 2007; Chen, 2009; Edwards and Susmel, 2001; Gray, 1996;

http://dx.doi.org/10.1016/j.econmod.2014.11.008
0264-9993/

Hamilton and Susmel, 1994; Henry, 2009; Klaassen, 2002; Wang and Theobald, 2008). However, these results are on occasion statistically insignificant or inconclusive (Dueker, 1997), and are less clear when assessed through some purely economic loss functions (Marcucci, 2005). Multivariate models of the conditional variance–covariance of returns, such as multivariate SWARCH (MV–SWARCH) and Markov switching BEK–GARCH, have also been used in order to compute time-varying optimal portfolio weights (Li, 2009) and hedge ratios (Lee and Yoder, 2007) or to assess contagion effects (Edwards and Susmel, 2001; Ramchand and Susmel, 1998).

Given the evidence of occasional discrete shifts in the conditional variance process, it is essential to test for the presence of multiple regimes in the conditional variance when a reasonable suspicion exists for structural change. In light of the calm and turbulence of the global and African stock markets during the 2002–2012 period, Markov-switching models may prove to be a more appropriate characterisation of stock return volatility than the popular single-regime GARCH. In contrast to the growing body of international literature, few studies of African financial market volatility incorporate regime-switching effects. Knedlik and Scheufele (2008) and Duncan and Liu (2009) both address the issue of forecasting South African currency crises within a SWARCH framework. Babikir et al. (2010) test the forecasting performance of an MS–GARCH model against a GJR-GARCH, using South African stock market data.

As far as could be determined, however, stock market volatility dynamics within and across multiple African markets have not been considered under a regime-switching framework in the literature. Given the importance of accounting for structural shifts in any time series analysis, it is important to contribute to this gap in the current body of knowledge. The current and future development of derivatives markets in the more financially sophisticated African markets (African Development Bank, 2010) will further increase the potential impact of such a study, as accurate stock return volatility estimates are required for the effective implementation of risk management strategies. The estimated volatility dynamics within and across African markets will also provide an indication of the nature and extent of contagion effects on the continent, with important implications for portfolio managers and policy makers alike.

The principal aim of this paper is to establish whether accounting for structural changes through the use of Markov-switching models improves estimates and forecasts of stock return volatility within and across selected African countries, namely South Africa, Kenya, Mauritius, Morocco and Nigeria. As such, the paper aims to address the following questions: (1) Do the univariate Markov-switching models of conditional variance provide a superior in-sample fit to the conventional single-state models? (2) Do Markov-switching models produce superior forecasts of the conditional variance to the conventional single-state models? (3) Do multivariate Markov-switching models of conditional covariance provide a more appropriate and accurate characterisation of stock return volatility and interdependence than the conventional single-state models? In addition to these goals, and as a by-product of the Markov-switching multivariate estimation, the paper aims to establish whether the correlation of returns between African markets changes across volatility states and the extent to which country-specific volatility states are dependent on one another.

The remainder of the paper is structured as follows: Section 2 provides an overview of the data and methods employed in this paper, including the relevant models and performance tests used. Section 3 presents the empirical results, while Section 4 concludes.

2. Data and econometric methods

2.1. Data

To accurately compare the properties of each market, it is appropriate to ensure that the return series are compiled according to a standardised method. For this reason, the data used are daily returns on the Morgan Stanley Capital International (MSCI) standard country indices for South Africa, Kenya, Nigeria, Mauritius and Morocco. All the series are obtained from *Thomson Reuters Datastream*. Although available, Egypt was excluded from the analysis in order to avoid the problem of missing data associated with the closing of the Egyptian stock market during the Arab Spring. Tunisia was also excluded despite availability, as observations for this market only begin in 2004. The series were selected so as to balance the need for adequate cross-continental representation with the need to analyse the longest possible history of returns. Despite the exclusion of Egypt and Tunisia, the series include a sufficiently diverse set of markets (both in terms of geographical dispersion and GDP growth) for the purposes of this paper. In addition, while African stock markets tend to respond to news from other African countries (Alagidede, 2010; Ntim, 2012), the overall market capitalisation, number of listed companies, and market capitalisation as percentage of GDP are generally among the highest across the five selected markets relative to other African countries (CMA, 2010; World Bank, 2014).

Since the daily return observations for Kenya, Nigeria and Mauritius are available from 3 June 2002, the South African and Moroccan series are also taken from this date. This is primarily because it is mathematically necessary to employ a uniform sample size when conducting the multivariate analysis. Also, a consistent sample size makes parameter estimates comparable across markets when conducting the univariate analysis. Thus, the sample for each series spans 3 June 2002 to 1 June 2012, encompassing 2610 daily observations. Daily returns (r) are calculated as:

$$r_t = 100 \times [\ln(P_t) - \ln(P_{t-1})] \tag{1}$$

which expresses daily returns in continuously compounded percentage terms. A graphical plot and the summary statistics of each series are presented in Fig. 1 and Table 1, respectively.

A few noteworthy patterns emerge from Table 1. The mean daily return figures range between the statistically insignificant 0.03% for Nigeria and 0.07% for Mauritius (corresponding to annual returns of roughly 7.15% and 17.6%, respectively). The markets with third lowest and lowest mean returns, South Africa and Nigeria, also exhibit the largest ex post variance of returns. South Africa, Mauritius and Morocco show significant degrees of skewness, whereas Kenya and Nigeria do not. Furthermore, all of the series exhibit statistically significant excess kurtosis, explaining the rejection of the null hypothesis of normally distributed returns in each case (as shown by the Jarque–Bera statistic). This is unsurprising, as the rejection of normally distributed returns is in line with much of the empirical literature using high frequency financial data (cf. Canarella and Pollard, 2007; Dueker, 1997).

Considering extreme values, the global financial crisis had a clear impact on all of the markets studied. Excluding Nigeria, every series experiences either a maximum or minimum (or both) observation during the month of October 2008. Another common high-volatility period seems to be in 2003, in which both Kenya and Nigeria exhibit extreme values. According to the Ljung–Box Q-statistics, all of the return series exhibit positive autocorrelation, which is a common finding within studies of emerging and frontier markets (Canarella and Pollard, 2007). Since the focus of this study is explicitly on the conditional variance and covariance of returns, consistent with existing work (Canarella and Pollard, 2007; Hamilton and Susmel, 1994; Lo and MacKinlay, 1990; Ramchand and Susmel, 1998) an AR(1) model is used to capture mean returns for all the countries considered. In addition to the autocorrelation found in the returns, each market exhibits positive autocorrelation in the squared returns, suggesting the presence of ARCH effects in the data. Interestingly, the South African market seems to be associated with the weakest autocorrelation in returns but the strongest autocorrelation in squared returns.

From Fig. 1, it appears that besides the occasional idiosyncratic shock, most series display a prolonged period of low volatility from