



ELSEVIER

Available online at www.sciencedirect.com

SCIENCE @ DIRECT®

JOURNAL OF
Econometrics

Journal of Econometrics 131 (2006) 179–215

www.elsevier.com/locate/econbase

Volatility comovement: a multifrequency approach

Laurent E. Calvet^{a,b,c,*}, Adlai J. Fisher^d, Samuel B. Thompson^a

^a*Department of Economics, Littauer Center, Harvard University, Cambridge, MA 02138, USA*

^b*HEC School of Management, Paris, France*

^c*NBER*

^d*Department of Finance, Sauder School of Business, University of British Columbia, 2053 Main Mall, Vancouver, BC, Canada V6T 1Z2*

Available online 7 March 2005

Abstract

We implement a multifrequency volatility decomposition of three exchange rates and show that components with similar durations are strongly correlated across series. This motivates a bivariate extension of the Markov-Switching Multifractal (MSM) introduced in Calvet and Fisher (J. Econ. 105 (2001) 27, J. Financ. Econ. 2 (2004) 49). Bivariate MSM is a stochastic volatility model with a closed-form likelihood. Estimation can proceed by maximum likelihood for state spaces of moderate size, and by simulated likelihood via a particle filter in high-dimensional cases. We estimate the model and confirm its main assumptions in likelihood ratio tests. Bivariate MSM compares favorably to a standard multivariate GARCH both in- and out-of-sample. A parsimonious multifrequency factor structure is finally proposed for multivariate settings with potentially many assets.

© 2005 Elsevier B.V. All rights reserved.

JEL classification: C13; C32; F37; G15

Keywords: Multivariate MSM; Maximum likelihood; Particle filter; Markov-switching; Stochastic volatility; Multifrequency volatility decomposition; Value-at-risk; Quantile forecasts

*Corresponding author. Department of Economics, Littauer Center, Harvard University, Cambridge, MA 02138, USA. Tel.: +1 617 495 4129; fax: +1 617 495 8570.

E-mail addresses: laurent_calvet@harvard.edu (L.E. Calvet), af@sauder.ubc.ca (A.J. Fisher), sthompson@harvard.edu (S.B. Thompson).

0304-4076/\$ - see front matter © 2005 Elsevier B.V. All rights reserved.

doi:10.1016/j.jeconom.2005.01.008

1. Introduction

A growing body of research investigates the comovement of volatility in financial series. The motivation underlying this effort is well-known. Joint movements in volatility influence the distribution of portfolio returns, and thus play an important role in risk management, portfolio selection, and derivative pricing. Comovements in volatility also help our understanding of financial markets, and shed light on issues such as contagion and the transmission of shocks through the financial system.¹ This motivation is particularly strong in the exchange rate literature, where first moments of currency returns relate weakly to fundamentals at medium frequencies, and movements in volatility can be large and persistent (e.g., Meese and Rogoff, 1983; Rogoff, 1999; Sarno and Taylor, 2002; Clarida et al., 2003).²

Multivariate GARCH, pioneered by Kraft and Engle (1982) and Bollerslev et al. (1988), is perhaps the most commonly used class of models. A natural extension of GARCH, these processes assume that a vector transform of the covariance matrix can be written as a linear combination of its lagged values and the return innovations. Andersen et al. (1999) show that these models perform well relative to competing alternatives. General formulations are, however, hampered by difficulties. The dimensionality of the parameter space grows quickly with the number of assets, and positive-definiteness of the covariance-matrix is not easily guaranteed. This has led to simplified versions including the constant-conditional correlation GARCH specification introduced by Bollerslev (1990).³ Multivariate stochastic volatility processes share these difficulties, although this has not prevented the growth of an impressive literature, including work by Harvey et al. (1994), Andersen et al. (2002), and Johannes et al. (2002).

We propose a new approach based on a recent advance in univariate time series, the Markov-switching multifractal (MSM) of Calvet and Fisher (2001, 2004). This earlier research uses Markov-switching to develop the first weakly convergent sequence of discrete filters for time-stationary multifractal diffusions. In MSM, total volatility is the multiplicative product of a large number of random components that are independent and statistically identical except for heterogeneity in their durations. The construction is parsimonious and delivers volatility persistence, substantial outliers, and a decomposition of volatility into frequency-specific components. MSM compares favorably to earlier specifications both in- and out-of-sample. Univariate multifractal forecasts slightly improve on GARCH(1,1) at daily and weekly intervals, and provide considerable gains in accuracy at horizons of 10–50 days.

This paper investigates comovement in MSM volatility components across exchange rates. We consider three series, the German Mark, the British Pound and the Japanese Yen, all vs. the US Dollar over the period 1973–2003. Our results show

¹See for instance Engle et al. (1990a), or Edwards and Susmel (2003) and the references therein.

²See Lyons (1995, 2001) for stronger evidence at high frequency.

³Researchers have also considered weaker restrictions (Engle and Kroner, 1995; Engle and Mezrich, 1996; Engle, 2002), factor structures (e.g., Engle, 1987; Diebold and Nerlove, 1989; Engle et al., 1990b), and estimation methods other than maximum likelihood (ML) (Ledoit et al., 2003).

that components from different series with similar frequencies tend to move together. In contrast, components with very different frequencies display less correlation, both within and across series. We attempt to relate MSM volatility components to macroeconomic indicators, and find no robust pattern using variables such as GDP, inflation, money supply, interest rates, and stock market volatility. On the other hand, oil and gold prices both correlate positively with currency volatility over the past three decades, consistent with the view that these commodities act as proxies for global economic and political risk. An MSM volatility decomposition reveals that these correlations exist only at low frequencies. This result encourages the econometrician to be cautious about the out-of-sample behavior of these apparent regularities.

Our findings motivate the construction of a bivariate model of currency volatility. This specification, called bivariate MSM, offers several advantages. It is relatively parsimonious, as the number of parameters is independent of \bar{k} . Positive semi-definiteness of the covariance matrix is trivially guaranteed. Furthermore, the likelihood function can be written in closed-form and ML estimation can be implemented for relatively small state spaces.

To accommodate a large number of frequencies, we develop a particle filter that permits convenient inference and forecasting using simulations. The good performance of the method is checked in Monte Carlo experiments. The algorithm broadens the range of computationally tractable models to include cases where the number of volatility components is very large, and to cases where the state variables are drawn from continuous rather than binomial distributions. This innovation thus opens econometric research on multifractal processes to a much wider range of specifications in both the univariate and multivariate cases.

We estimate the bivariate model by ML and verify that the goodness of fit increases with the number of frequency-specific volatility components. Likelihood ratio (LR) tests confirm that the main assumptions of the model are empirically valid. Bivariate MSM compares favorably to constant correlation GARCH (CC-GARCH) in-sample. Out-of-sample, integral transforms and the Cramer–von Mises (CVM) statistic indicate that, in contrast to CC-GARCH, bivariate MSM captures well the conditional distribution of a variety of currency portfolios. MSM also provides reasonable measures of value-at-risk (VaR), while CC-GARCH tends to underestimate the riskiness of a currency position.

We conclude the paper with generalizations of the model to a larger number of assets. We show that bivariate MSM easily extends to larger economies, but can become complicated in its general formulation. We thus propose a factor model of multifrequency stochastic volatility, specified by a number of volatility parameters that grows linearly in the number of assets. Estimation can be conducted by maximizing the closed-form likelihood or by implementing the particle filter.

The rest of the paper is organized as follows. Section 2 reviews univariate MSM and relates the volatility components to other financial and macroeconomic indicators. Section 3 introduces the bivariate model. Section 4 develops inference methods, including likelihood estimation and the particle filter. Empirical results are

reported in Section 5. Extensions to many assets are discussed in Section 6. Unless stated otherwise, all proofs are given in the Appendix A.

2. Comovement of univariate volatility components

2.1. The univariate stochastic volatility model

We begin by reviewing the Markov-switching multifractal (MSM), a discrete-time Markov process with multi-frequency stochastic volatility. Consider an economic series p_t defined in discrete time on the regular grid $t = 0, 1, 2, \dots, \infty$. In applications, p_t is the log-price of a financial asset or exchange rate. We consider an economy with \bar{k} volatility components $M_{1,t}, M_{2,t}, \dots, M_{\bar{k},t}$, which decay at heterogeneous frequencies $\gamma_1, \dots, \gamma_{\bar{k}}$. The notation MSM (\bar{k}) refers to versions of the model with \bar{k} frequencies.

The innovations $x_t \equiv p_t - p_{t-1}$ are specified as

$$x_t = (M_{1,t}M_{2,t} \dots M_{\bar{k},t})^{1/2}\varepsilon_t, \tag{2.1}$$

where the random variables ε_t are IID standard Gaussians $\mathcal{N}(0, \sigma^2)$. The random *multipliers* or *volatility components* $M_{k,t}$ are persistent, non-negative and satisfy $\mathbb{E}(M_{k,t}) = 1$. We assume for simplicity that the multipliers $M_{1,t}, M_{2,t} \dots M_{\bar{k},t}$ at a given time t are statistically independent. The parameter σ then equals the unconditional standard deviation of the innovation x_t .

Eq. (2.1) defines a return process with stochastic volatility $\sigma_t = \sigma(M_{1,t}M_{2,t} \dots M_{\bar{k},t})^{1/2}$. We conveniently stack the period t volatility components into the $1 \times \bar{k}$ row vector

$$M_t = (M_{1,t}, M_{2,t}, \dots, M_{\bar{k},t}).$$

The vector M_t is assumed to be first-order Markov, and is thus called the *volatility state*. The econometrician observes the returns x_t but not the vector M_t itself. MSM is thus a hidden Markov chain model of volatility. The latent state M_t is inferred recursively by Bayesian updating, and estimation is possible by ML. MSM is thus a tractable high-dimensional version of the regime-switching models advocated by Hamilton (1989).

The dynamics of the various volatility components $M_{k,t}$ are identical except for time scale. Assume that the state vector M_{t-1} has been constructed in date $t - 1$. For each $k \in \{1, \dots, \bar{k}\}$, the next period multiplier $M_{k,t}$ is either drawn from a fixed distribution M with probability γ_k , or otherwise remains equal to its current value: $M_{k,t} = M_{k,t-1}$. The dynamics of $M_{k,t}$ can be summarized as

$$\begin{aligned} M_{k,t} &\text{ drawn from distribution } M && \text{ with probability } \gamma_k, \\ M_{k,t} &= M_{k,t-1} && \text{ with probability } 1 - \gamma_k. \end{aligned}$$

The switching events and new draws from M are assumed to be independent across k and t . The volatility components $M_{k,t}$ thus differ in their transition probabilities γ_k but not in their marginal distribution M .

The transition probabilities are specified as $\gamma_k = 1 - (1 - \gamma_{\bar{k}})^{(b^{k-\bar{k}})}$, where $\gamma_{\bar{k}} \in (0, 1)$ and $b \in (1, \infty)$. This specification is introduced in Calvet and Fisher (2001) in connection with the discretization of a Poisson arrival process. The pair $(\gamma_{\bar{k}}, b)$ thus provides a numerically convenient specification for the transition probabilities.

The multifractal construction imposes only minimal restrictions on the marginal distribution of the multipliers: $M \geq 0$ and $\mathbb{E}M = 1$. While flexible parametric or even non-parametric specifications of M can be used, this paper focuses on the parsimonious setup in which M is drawn from a binomial random variable taking values $m_0 \in [1, 2]$ or $2 - m_0 \in [0, 1]$ with equal probability. The full parameter vector is then

$$\psi \equiv (m_0, \sigma, b, \gamma_{\bar{k}}) \in \mathbb{R}_+^4,$$

where m_0 characterizes the distribution of the multipliers, σ is the unconditional standard deviation of returns, and b and $\gamma_{\bar{k}}$ define the set of switching probabilities.

The multiplicative structure (2.1) is appealing to model the outliers and volatility persistence exhibited by financial time series. Changes in low-level multipliers lead to discrete shifts in volatility that can be maintained over long periods of time. Such risks are important, for example, to a long-term investor or to a market maker pricing long-lived options. In addition, variations in high-frequency multipliers help capture extreme tail events in short-run returns. This has obvious implications for pricing shorter-lived options or for calculating VaR.

In exchange rate series, the estimated duration of the most persistent component, $1/\gamma_1$, is typically of the same order as the length of the data. The process thus generates volatility cycles with periods proportional to the sample size, a property also apparent in the sample paths of long-memory processes. Fractionally integrated processes generate such patterns by assuming that an innovation linearly affects future periods at a hyperbolically declining weight. As a result, fractional integration tends to produce smooth volatility processes. By contrast, our approach generates long cycles with a switching mechanism that also gives abrupt volatility changes.⁴ The combination of long-memory behavior with sudden volatility movements in MSM has a natural appeal for financial econometrics.

The continuous-time version of MSM is conveniently constructed and lies outside the class of Itô diffusions when $\bar{k} \rightarrow \infty$. The sample paths are continuous but exhibit a high degree of heterogeneity in local behavior, which is characterized by a continuum of local Hölder exponents in any finite time interval. Calvet and Fisher (2001, 2002) fully develop the continuous-time limit.

⁴Long memory is often defined by a hyperbolic decline in the autocovariance function as the lag goes to infinity. As shown in Calvet and Fisher (2004), MSM mimics the hyperbolic autocorrelation in the size of the returns exhibited by many financial series (e.g., Ding et al., 1993). The multifractal model thus illustrates the difficulty of distinguishing between long memory and structural change in finite samples, as in Hidalgo and Robinson (1996) or Diebold and Inoue (2001).

2.2. Comovement of exchange rate volatility

The empirical analysis investigates daily returns on the Deutsche Mark (DM), Japanese Yen (JA) and British Pound (UK), all against the US Dollar. The returns are imputed from noon daily prices reported by the Federal Reserve Bank of New York.⁵ The series begin on 1 June 1973, shortly after the demise of the Bretton Woods fixed exchange rate system. The Deutsche Mark is replaced by the Euro at the beginning of 1999. Each series ends on 30 October 2003 and contains 7635 observations.

For each currency, we estimate MSM by ML on the entire sample and report the results in Table 1. The columns correspond to the number of frequencies \bar{k} varying from 1 to 8. The first column is thus a standard regime-switching process with only two possible levels of volatility. As \bar{k} increases, the number of states increases at the rate $2^{\bar{k}}$. The multiplier value m_0 tends to decline with \bar{k} . With a larger number of frequencies, less variability is required from each individual component to match the volatility fluctuations of the data. Estimates of σ fluctuate across \bar{k} with no apparent trend. When $\bar{k} = 1$, the parameter $\hat{\gamma}_{\bar{k}}$ indicates that the single multiplier has a duration of a few weeks. As \bar{k} increases, the switching probability of the highest frequency multiplier increases until a switch occurs almost every day for large \bar{k} . At the same time, the growth rate \hat{b} decreases steadily with \bar{k} . In the DM series with $\bar{k} = 8$, a switch in the lowest frequency multiplier occurs approximately once every 8 years, or about one fourth the sample size. Thus, as \bar{k} increases, frequencies tend to span a wider range while becoming more tightly spaced.

We use the ML estimates to compute, for each currency, the smoothed state probabilities (Kim, 1993) and the expectation of the multipliers conditional on the entire sample: $\hat{M}_{k,t} = \mathbb{E}(M_{k,t} | x_1, \dots, x_T)$. The correlations of the smoothed components $\hat{M}_{k,t}$ are reported in Table 2. In the first panel, we see that different components of the DM exchange rate are moderately correlated, and correlation decreases in the distance between frequencies.⁶ We report only DM results for space constraints, but obtain similar results with the UK and JA series. The second and third panels of Table 2 show inferred comovement of the DM components with JA and UK. Correlation between the smoothed beliefs $\hat{M}_{k,t}^{\alpha}$ and $\hat{M}_{k',t}^{\beta}$ of two currencies tends to be high when k and k' are close, and is low otherwise. This suggests that the volatility components of two exchange rates are correlated only if their frequencies are similar.

The interpretation is slightly complicated, however, by the fact that the set of volatility frequencies is not identical across currencies (Table 1). To address this issue, we now introduce a simple bivariate model in which currencies are statistically independent but have identical frequency parameters b and $\gamma_{\bar{k}}$. The log-likelihood of

⁵More specifically, the data consist of buying rates for wire transfers at 12:00 PM Eastern time.

⁶Note that since the econometrician does not directly observe the multipliers, some degree of correlation in smoothed beliefs is consistent with the independence of the unobserved components $M_{k,t}$ and $M_{k',t}$, $k \neq k'$.

Table 1
Univariate MLE

	$\bar{k} = 1$	2	3	4	5	6	7	8
<i>Deutsche Mark</i>								
\hat{m}_0	1.617 (0.019)	1.556 (0.015)	1.535 (0.012)	1.472 (0.012)	1.445 (0.013)	1.396 (0.012)	1.365 (0.011)	1.338 (0.011)
$\hat{\sigma}$	0.672 (0.012)	0.649 (0.017)	0.594 (0.013)	0.567 (0.015)	0.504 (0.016)	0.537 (0.027)	0.549 (0.020)	0.552 (0.021)
$\hat{\gamma}_{\bar{k}}$	0.074 (0.002)	0.086 (0.018)	0.841 (0.096)	0.779 (0.082)	0.812 (0.083)	0.909 (0.103)	0.979 (0.036)	0.998 (0.008)
\hat{b}	—	6.85 (2.44)	34.31 (10.55)	11.86 (1.99)	9.02 (1.24)	5.83 (0.82)	4.67 (0.60)	3.82 (0.49)
$\ln L$	-7121.92	-6975.92	-6916.81	-6900.06	-6891.67	-6888.91	-6885.60	-6885.90
<i>Japanese Yen</i>								
\hat{m}_0	1.783 (0.011)	1.774 (0.009)	1.688 (0.011)	1.644 (0.011)	1.579 (0.010)	1.567 (0.010)	1.559 (0.010)	1.508 (0.010)
$\hat{\sigma}$	0.632 (0.011)	0.537 (0.009)	0.568 (0.019)	0.473 (0.017)	0.473 (0.023)	0.634 (0.023)	0.514 (0.019)	0.508 (0.017)
$\hat{\gamma}_{\bar{k}}$	0.208 (0.022)	0.358 (0.038)	0.276 (0.048)	0.713 (0.082)	0.861 (0.053)	0.894 (0.060)	0.894 (0.058)	0.977 (0.030)
\hat{b}	—	147.47 (59.61)	11.76 (2.02)	15.73 (2.67)	9.13 (1.18)	8.22 (0.99)	7.60 (0.87)	5.88 (0.74)
$\ln L$	-6776.19	-6421.01	-6279.02	-6216.85	-6196.55	-6184.90	-6181.29	-6174.96
<i>British Pound</i>								
\hat{m}_0	1.708 (0.013)	1.666 (0.013)	1.640 (0.011)	1.612 (0.014)	1.574 (0.011)	1.529 (0.012)	1.498 (0.011)	1.457 (0.010)
$\hat{\sigma}$	0.606 (0.009)	0.580 (0.018)	0.523 (0.018)	0.516 (0.016)	0.431 (0.015)	0.455 (0.017)	0.385 (0.013)	0.380 (0.014)
$\hat{\gamma}_{\bar{k}}$	0.113 (0.016)	0.213 (0.036)	0.271 (0.065)	0.549 (0.086)	0.617 (0.074)	0.782 (0.078)	0.817 (0.083)	0.959 (0.001)
\hat{b}	—	18.69 (4.84)	13.92 (2.68)	14.39 (2.67)	11.59 (1.84)	8.49 (1.16)	6.83 (0.87)	5.33 (0.04)
$\ln L$	-6220.55	-5987.37	-5882.60	-5826.92	-5792.97	-5778.58	-5771.92	-5770.20

This table shows ML estimation results for binomial MSM. Columns correspond to the number \bar{k} of volatility components. Asymptotic standard errors are in parenthesis.

the two series is then

$$\ln L(x_t^\alpha; m_0^\alpha, \sigma_\alpha, b, \gamma_{\bar{k}}) + \ln L(x_t^\beta; m_0^\beta, \sigma_\beta, b, \gamma_{\bar{k}}), \quad (2.2)$$

where L denotes the likelihood of univariate MSM. This specification, called the *combined univariate*, is an important building block of the bivariate model introduced in the next section.

In Table 3, we report empirical results for the combined univariate. Panel A shows ML estimates for the (DM, JA) series. Some parameter estimates differ noticeably from the unrestricted univariate results in Table 1, but substantial discrepancies only

Table 2
Correlation of smoothed univariate volatility components

	DM1	DM2	DM3	DM4	DM5	DM6	DM7	DM8	$ x_{DM} $	x_{DM^2}
DM1	1.000	0.762	0.377	0.174	0.099	0.066	0.031	0.020	0.189	0.115
DM2	0.762	1.000	0.600	0.328	0.151	0.093	0.043	0.028	0.255	0.174
DM3	0.377	0.600	1.000	0.603	0.307	0.168	0.077	0.052	0.312	0.245
DM4	0.174	0.328	0.603	1.000	0.738	0.432	0.201	0.137	0.374	0.295
DM5	0.099	0.151	0.307	0.738	1.000	0.792	0.420	0.297	0.463	0.373
DM6	0.066	0.093	0.168	0.432	0.792	1.000	0.770	0.610	0.667	0.539
DM7	0.031	0.043	0.077	0.201	0.420	0.770	1.000	0.961	0.887	0.713
DM8	0.020	0.028	0.052	0.137	0.297	0.610	0.961	1.000	0.894	0.716
$ x_{DM} $	0.189	0.255	0.312	0.374	0.463	0.667	0.887	0.894	1.000	0.872
x_{DM^2}	0.115	0.174	0.245	0.295	0.373	0.539	0.713	0.716	0.872	1.000
JA1	0.590	0.287	0.036	0.020	0.036	0.032	0.012	0.007	0.051	0.002
JA2	0.611	0.302	0.048	0.023	0.038	0.034	0.013	0.008	0.056	0.006
JA3	0.788	0.440	0.172	0.063	0.065	0.048	0.021	0.013	0.104	0.048
JA4	0.368	0.185	0.162	0.064	0.030	0.036	0.020	0.013	0.073	0.062
JA5	0.157	0.177	0.150	0.231	0.169	0.109	0.053	0.036	0.103	0.084
JA6	0.058	0.062	0.127	0.279	0.349	0.284	0.155	0.111	0.192	0.174
JA7	0.029	0.023	0.032	0.106	0.206	0.321	0.312	0.267	0.284	0.258
JA8	0.012	0.008	0.011	0.043	0.095	0.209	0.339	0.353	0.333	0.303
$ x_{JA} $	0.187	0.108	0.092	0.134	0.177	0.256	0.328	0.327	0.363	0.346
x_{JA^2}	0.091	0.048	0.065	0.101	0.142	0.209	0.261	0.256	0.297	0.344
UK1	0.819	0.525	0.170	0.081	0.052	0.042	0.018	0.011	0.114	0.054
UK2	0.730	0.558	0.246	0.165	0.094	0.062	0.028	0.018	0.134	0.073
UK3	0.464	0.526	0.254	0.163	0.072	0.050	0.023	0.015	0.128	0.086
UK4	0.251	0.505	0.308	0.195	0.075	0.049	0.021	0.014	0.143	0.111
UK5	0.070	0.131	0.516	0.571	0.365	0.200	0.091	0.062	0.213	0.169
UK6	0.149	0.162	0.239	0.440	0.536	0.423	0.228	0.162	0.291	0.242
UK7	0.082	0.079	0.092	0.185	0.319	0.463	0.431	0.362	0.407	0.354
UK8	0.030	0.030	0.035	0.074	0.145	0.301	0.473	0.488	0.471	0.409
$ x_{UK} $	0.168	0.213	0.221	0.254	0.273	0.360	0.462	0.462	0.564	0.524
x_{UK^2}	0.081	0.135	0.178	0.213	0.234	0.305	0.390	0.391	0.508	0.571

This table shows correlations from a frequency decomposition of binomial MSM with $\bar{k} = 8$ components for the univariate DM series with itself, JA, and UK. For each series, the smoothed components $\hat{M}_{k,t} = \mathbb{E}(M_{k,t}|x_1, \dots, x_T)$ of different volatility states are calculated. For convenience, we denote these probabilities by DM1, ..., DM8, JA1, ..., JA8, UK1, ..., UK8. The table then shows correlations of the DM decomposition with decompositions from all three series. Correlations are generally strongest near the diagonal.

seem to occur for low values of \bar{k} . For instance with $\bar{k} = 8$, we compare the log-likelihood of -13063.11 with the sum of unrestricted log-likelihoods, i.e. $-6885.90 - 6174.96 = -13060.86$, finding a difference of 2.25. Under the combined univariate, the LR statistic, 2×2.25 , is asymptotically distributed as a chi-squared with two degrees of freedom. This test statistic is not significant at any conventional level, confirming that the frequency restrictions are reasonable. The second part of Panel A

Table 3
Combined univariate results

	$\bar{k} = 1$	2	3	4	5	6	7	8		
<i>A. MLE estimation</i>										
DM–JA parameter estimates										
\hat{m}_0^{DM}	1.643 (0.013)	1.618 (0.014)	1.515 (0.013)	1.474 (0.013)	1.445 (0.011)	1.405 (0.013)	1.397 (0.012)	1.367 (0.011)		
\hat{m}_0^{JA}	1.775 (0.018)	1.757 (0.010)	1.687 (0.010)	1.638 (0.011)	1.578 (0.010)	1.565 (0.010)	1.522 (0.010)	1.488 (0.014)		
$\hat{\sigma}_{DM}$	0.669 (0.010)	0.577 (0.011)	0.597 (0.015)	0.569 (0.018)	0.504 (0.015)	0.565 (0.021)	0.449 (0.018)	0.472 (0.018)		
$\hat{\sigma}_{JA}$	0.613 (0.016)	0.544 (0.010)	0.565 (0.018)	0.487 (0.018)	0.476 (0.023)	0.632 (0.024)	0.384 (0.016)	0.532 (0.027)		
$\hat{\gamma}_{\bar{k}}$	0.129 (0.013)	0.257 (0.026)	0.301 (0.068)	0.756 (0.057)	0.844 (0.055)	0.872 (0.054)	0.959 (0.036)	0.982 (0.022)		
\hat{b}	—	69.57 (21.57)	11.97 (2.04)	13.21 (1.61)	9.14 (0.95)	7.16 (0.65)	6.16 (0.57)	4.93 (0.45)		
$\ln L$	-13913.86	-13424.24	-13203.13	-13119.07	-13088.39	-13077.02	-13072.22	-13063.11		
LR tests against unrestricted univariate										
DM–JA	0.000	0.000	0.001	0.116	0.849	0.040	0.005	0.106		
DM–UK	0.136	0.002	0.000	0.037	0.004	0.004	0.004	0.111		
JA–UK	0.004	0.000	0.701	0.361	0.017	0.237	0.754	0.863		
<i>B. Correlation of smoothed volatility component beliefs</i>										
	DM1	DM2	DM3	DM4	DM5	DM6	DM7	DM8	$ x_{DM} $	x_{DM^2}
JA1	0.628	0.714	0.349	0.072	0.009	0.038	0.020	0.007	0.081	0.028
JA2	0.690	0.774	0.405	0.135	0.016	0.048	0.027	0.011	0.101	0.043
JA3	0.595	0.686	0.228	0.140	0.018	0.049	0.028	0.012	0.078	0.033
JA4	0.306	0.234	0.147	0.114	0.036	0.022	0.027	0.013	0.065	0.065
JA5	-0.019	-0.034	0.052	0.113	0.302	0.227	0.116	0.056	0.117	0.102
JA6	0.028	0.023	0.040	0.084	0.255	0.352	0.258	0.145	0.206	0.186
JA7	0.008	0.009	0.007	0.021	0.103	0.224	0.342	0.294	0.299	0.274
JA8	0.004	0.004	0.002	0.009	0.048	0.123	0.287	0.353	0.333	0.304
$ x_{JA} $	0.177	0.191	0.093	0.088	0.128	0.193	0.301	0.326	0.363	0.346
x_{JA^2}	0.087	0.094	0.039	0.066	0.100	0.156	0.243	0.254	0.297	0.344

Panel A shows maximum likelihood estimation results for the combined univariate model, which for two series α and β has likelihood $L(x_t^\alpha; m_0^\alpha, \sigma_\alpha, b, \gamma_{\bar{k}})L(x_t^\beta; m_0^\beta, \sigma_\beta, b, \gamma_{\bar{k}})$. This corresponds to the likelihood of two statistically independent univariate MSM processes constrained to have the same frequency parameters b and $\gamma_{\bar{k}}$. Columns correspond to the number of frequencies \bar{k} in the estimated model, and estimation results with asymptotic standard errors in parentheses are presented for the DM–JA currency pair only. The second part of Panel A shows p -values from a likelihood ratio test of the combined univariate against two unrestricted independent MSM processes. A low p -value indicates rejection of the restriction that frequency parameters are identical across currencies. Panel B then shows correlations from a frequency decomposition of the DM–JA combined univariate model with eight components. For each series, the smoothed probabilities $M_{k,t} = \mathbb{E}(M_{k,t}|x_1, \dots, x_T)$ of volatility states are calculated. For convenience, we denote these probabilities by DM1, ..., DM8, JA1, ..., JA8.

reports more generally the p -values of the LR tests corresponding to each frequency and currency combination. When $\bar{k} = 8$, the frequency restrictions are not rejected for any currency combination. Panel B then shows correlations between smoothed

volatility components for the (DM, JA) series under the combined univariate model. With frequencies now identical across currencies, we expect results to be stronger than in Table 2, and this is confirmed. Results for other currency pairs are similar. We thus find that: (i) restricting frequencies to be identical across currencies is reasonable, and (ii) components of similar frequencies tend to move together while components with very different frequencies do not. These conclusions are useful in developing a bivariate MSM specification in Section 3.

2.3. Currency volatility and macroeconomic indicators

We now investigate whether currency volatility comovement relates to other macroeconomic and financial variables. Earlier research leads us to be relatively pessimistic. For instance, the first moments of exchange rates are weakly linked to fundamentals (e.g., Meese and Rogoff, 1983; Andersen and Bollerslev, 1998; Rogoff, 1999; Sarno and Taylor, 2002). Variances are also difficult to explain, at least in stock market data (e.g., Schwert, 1989). We examine whether the new MSM multifrequency decomposition confirms these negative results.

We first consider IMF monthly data for 1973–2000, including monetary aggregates (M1, M2 and M3), short and long interest rates, producer price index, consumer price index, wages and the growth rate of industrial production. We compute the correlation between monthly volatility and the macro variables of each country, their difference and the absolute value of their difference. We use several measures of volatility, such as the absolute value of the monthly return, realized monthly volatility, and MSM volatility components. In unreported work using a variety of lag structures, we find no robust link between currency volatility and these variables. These results are consistent with the findings of Andersen and Bollerslev (1998), who show that macro announcements induce volatility shocks that are of comparable magnitude to daily volatility. It is thus not surprising that little impact is found at the monthly frequency.

Economic theory suggests that exchange rates might be more strongly linked to equity markets, since both classes of instruments incorporate forward-looking information about rates of return, national economic conditions and corporate profits.⁷ In Table 4, we investigate the comovement of each currency with volatility in domestic and US stock markets. Daily returns are imputed from the US value-weighted CRSP index, the German CDAX Composite Price Index, the UK FT-Actuaries All Share Index and the Japanese Nikkei 225 Stock Average. Realized monthly stock volatility (RV) is computed as the sum of squared daily returns. We compare it with the currency return, absolute return, RV and MSM frequency-specific components. The reported correlation is positive for the Deutsche Mark and the Yen, and weakly negative for the Pound. We thus find no robust link between currency and stock volatility.

Oil prices are often viewed as proxies for global economic and political uncertainty (e.g., Hamilton, 2003). As seen in Table 4, the dollar price of oil correlates positively

⁷See Sarno and Taylor (2002) for a recent review of the economics of exchange rates.

Table 4
Correlation of exchange rates with other financial prices

	x_t	$ x_t $	RV	MSM volatility component beliefs							
				$k = 1$	2	3	4	5	6	7	8
<i>Deutsche Mark</i>											
CRSP RV	-0.0849	0.1712	0.0904	0.0930	0.0360	0.0316	0.1964	0.1421	0.0678	-0.0056	-0.0479
DAX RV	-0.1515	0.0916	0.1205	0.2087	0.1296	0.1086	0.1541	0.1041	0.0111	-0.0363	-0.0395
Oil	0.1142	0.0235	0.2138	0.5254	0.4859	0.2685	0.1649	0.0241	-0.0178	-0.1062	-0.1137
Gold	0.0774	0.0652	0.1642	0.7378	0.4979	0.1334	0.0571	-0.0104	-0.0613	-0.0755	-0.0497
<i>Japanese Yen</i>											
CRSP RV	-0.0969	0.1696	0.2097	-0.1841	-0.1924	-0.2137	-0.1439	-0.0920	0.0011	-0.0136	0.0221
NIKKEI RV	-0.0280	0.1155	0.2291	0.3069	0.3013	0.2764	-0.3071	-0.0965	-0.0526	-0.0781	-0.0400
Oil	0.0465	0.0429	0.0046	0.2803	0.2858	0.2838	0.1306	-0.1608	0.0335	-0.0422	0.0052
Gold	-0.0135	0.1323	0.1444	0.5233	0.5306	0.5332	0.2071	0.0586	0.0417	-0.0633	-0.1053
<i>British Pound</i>											
US CRSP RV	0.0480	0.0590	-0.0746	-0.0299	0.0907	0.0496	-0.0597	0.0291	0.0570	-0.0342	-0.0405
FTSE RV	0.0113	0.0390	-0.0986	-0.4642	-0.2036	-0.0865	-0.1638	0.1450	0.0405	-0.0329	-0.0518
Oil	-0.0900	0.0864	0.2583	0.5495	0.3972	0.2652	0.2846	0.2087	0.0766	-0.0175	-0.0455
Gold	0.0396	0.0837	0.1466	0.6963	0.5087	0.4084	0.2824	-0.1019	0.0658	0.0083	-0.0243

This table investigates for each country the comovement between exchange rates and four financial variables: the monthly RV on the US and domestic stock market, the oil price (in USD/barrel) and the gold price (in USD/oz). Monthly realized volatilities are imputed as the sum of squared daily returns. Correlation between currency and equity RV is positive for DM and JA, but negative for UK. Oil and gold prices are positively correlated to exchange rate volatility for all countries, and the MSM decomposition reveals that this result is primarily a low-frequency phenomenon.

with the RV of DM and UK,⁸ and the MSM decomposition further reveals that this is primarily a low-frequency phenomenon. The results become more intriguing for JA. While the raw oil price shows little correlation with the RV of the Yen, it is again strongly correlated with low-frequency MSM components. The MSM decomposition thus finds evidence of a regularity that direct analysis of RV would not uncover. Similar results are obtained with gold, further suggesting that currency volatility and certain commodity prices may be linked at low frequencies through an unidentified global risk factor. As in [Stock and Watson \(2003\)](#), we view these findings with caution since it is unclear whether oil and gold prices will continue to be effective proxies for global risk in the future.

Our results thus deepen the puzzle regarding the link between exchange rates and fundamentals. Volatility components are strongly correlated across currencies but only weakly related to other macroeconomic and financial variables. This motivates the development of a multivariate multifrequency model of exchange rates.

⁸We use the domestic first purchase price of crude oil expressed in dollars per barrel provided by Global Insight/DRI.

3. A bivariate multifrequency model

3.1. The stochastic volatility specification

We consider two financial series α and β defined on the regular grid $t = 0, 1, 2, \dots, \infty$. Their log-returns⁹ x_t^α and x_t^β in period t are stacked into the column vector

$$x_t = \begin{bmatrix} x_t^\alpha \\ x_t^\beta \end{bmatrix} \in \mathbb{R}^2.$$

As in univariate MSM, volatility is stochastic and hit by shocks of heterogeneous frequencies indexed by $k \in \{1, \dots, \bar{k}\}$. For every frequency k , the currencies have volatility components $M_{k,t}^\alpha$ and $M_{k,t}^\beta$. Consider

$$M_{k,t} = \begin{bmatrix} M_{k,t}^\alpha \\ M_{k,t}^\beta \end{bmatrix} \in \mathbb{R}_+^2.$$

The period- t volatility column vectors $M_{k,t}$ are stacked into the $2 \times \bar{k}$ matrix

$$M_t = (M_{1,t}; M_{2,t}; \dots; M_{\bar{k},t}).$$

Each row of the matrix M_t contains the volatility components of a particular currency, while each column corresponds to a particular frequency. As in univariate MSM, we assume that $M_{1,t}, M_{2,t}, \dots, M_{\bar{k},t}$ at a given time t are statistically independent. The main task is to choose appropriate dynamics for each vector $M_{k,t}$.

Economic intuition suggests that volatility arrivals are correlated but not necessarily simultaneous across currency markets. For this reason, we allow arrivals across series to be characterized by a correlation coefficient λ . Assume that the volatility vector $M_{k,s}$ has been constructed up to date $t - 1$. In period t , each series $c \in \{\alpha, \beta\}$ is hit by an arrival with probability γ_k . Let $1_{k,t}^c$ denote the random variable equal to 1 if there is an arrival on $M_{k,t}^c$, and equal to 0 otherwise. The arrival vector $1_{k,t} = (1_{k,t}^\alpha, 1_{k,t}^\beta)$ is specified to be IID, and its unconditional distribution is defined by the following three conditions. First, the arrival vector is symmetrically distributed: $(1_{k,t}^\beta, 1_{k,t}^\alpha) \stackrel{d}{=} (1_{k,t}^\alpha, 1_{k,t}^\beta)$. Second, the switching probability of a series is equal to an exogenous constant:

$$\mathbb{P}(1_{k,t}^\alpha = 1) = \gamma_k.$$

Third, there exists $\lambda \in [0, 1]$ such that

$$\mathbb{P}(1_{k,t}^\beta = 1 | 1_{k,t}^\alpha = 1) = (1 - \lambda)\gamma_k + \lambda.$$

As shown in the Appendix A, these three conditions define a unique distribution for $1_{k,t}$. Arrivals are independent if $\lambda = 0$ and simultaneous if $\lambda = 1$. We easily check that λ is the unconditional correlation between $1_{k,t}^\alpha$ and $1_{k,t}^\beta$.

⁹If X_t^α denote the value of the exchange rate at date t , the log-return is $x_t^\alpha = \ln(X_t^\alpha / X_{t-1}^\alpha)$.

Given the realization of the arrival vector $1_{k,t}$, the construction of the volatility components $M_{k,t}$ is based on a bivariate distribution $M = (M^\alpha, M^\beta) \in \mathbb{R}_+^2$. We briefly postpone the choice of M , and assume for now that it is defined by two parameters m_0^α and m_0^β . If arrivals hit both series ($1_{k,t}^\alpha = 1_{k,t}^\beta = 1$), the state vector $M_{k,t}$ is drawn from M . If only series $c \in \{\alpha, \beta\}$ receives an arrival, the new component $M_{k,t}^c$ is sampled from the marginal M^c of the bivariate distribution M . Finally, $M_{k,t} = M_{k,t-1}$ if there is no arrival ($1_{k,t}^\alpha = 1_{k,t}^\beta = 0$). The construction thus implies that arrival vectors and draws from M are independent across k and t . In vector notation, it can be summarized as

$$M_{k,t} \stackrel{d}{=} M_{k,t-1} + 1_{k,t} * (M - M_{k,t-1}),$$

where $*$ denote element by element multiplication. The volatility components $M_{k,t}$ differ in their transition probabilities γ_k , but not in their marginal distribution M or arrival correlation λ . These features greatly contribute to the parsimony of the model.

The volatility vectors $M_{k,t}$ are persistent, non-negative and satisfy $\mathbb{E}(M_{k,t}) = \mathbf{1}$, where $\mathbf{1} = (1, 1)'$. Consistent with previous notation, let $g(M_t)$ denote the 2×1 vector $M_{1,t} * M_{2,t} * \dots * M_{\bar{k},t}$. We specify the return vector x_t as

$$x_t = [g(M_t)]^{1/2} * \varepsilon_t, \tag{3.1}$$

where the column vectors $\varepsilon_t \in \mathbb{R}^2$ are IID Gaussian $\mathcal{N}(0, \Sigma)$. The covariance matrix Σ is conveniently written in the standard form

$$\Sigma = \begin{bmatrix} \sigma_\alpha^2 & \rho_\varepsilon \sigma_\alpha \sigma_\beta \\ \rho_\varepsilon \sigma_\alpha \sigma_\beta & \sigma_\beta^2 \end{bmatrix}.$$

The construction thus permits correlation in volatility across series through the bivariate distribution M , and correlation in returns through the Gaussian vector ε_t . As in the univariate case, the transition probabilities $(\gamma_1, \gamma_2, \dots, \gamma_{\bar{k}})$ are defined as

$$\gamma_k = 1 - (1 - \gamma_{\bar{k}})^{(b^{k-\bar{k}})}, \tag{3.2}$$

where $\gamma_{\bar{k}} \in (0, 1)$ and $b \in (1, \infty)$. This completes the specification of *bivariate MSM*.

We observe that under this model, individual returns satisfy

$$\begin{aligned} x_t^\alpha &= (M_{1,t}^\alpha M_{2,t}^\alpha \dots M_{k,t}^\alpha)^{1/2} \varepsilon_t^\alpha, \\ x_t^\beta &= (M_{1,t}^\beta M_{2,t}^\beta \dots M_{k,t}^\beta)^{1/2} \varepsilon_t^\beta. \end{aligned}$$

Their univariate dynamics thus coincide with univariate MSM. In particular, the parameter σ_c is again the unconditional standard deviation of each univariate series $c \in \{\alpha, \beta\}$. Bivariate MSM is overall specified by the eight parameters

$$\psi \equiv (\sigma_\alpha, \sigma_\beta, m_0^\alpha, m_0^\beta, b, \gamma_{\bar{k}}, \rho_\varepsilon, \lambda),$$

where σ_α and σ_β are the unconditional standard deviations of the return series, m_0^α and m_0^β determine the distribution of volatility components, $\gamma_{\bar{k}}$ their transition

probabilities, ρ_ε the correlation of the Gaussian innovations, and λ the correlation of arrivals across series.

The bivariate construction imposes few restrictions on the distribution of vector M . In the empirical applications of this paper, we investigate a simple specification where each M_{kt} is drawn from a bivariate binomial distribution $M = (M^\alpha, M^\beta)'$. The first element M^α takes values $m_0^\alpha \in [1, 2]$ and $m_1^\alpha = 2 - m_0^\alpha \in [0, 1]$ with equal probability. Similarly, M^β is either $m_0^\beta \in [1, 2]$ or $m_1^\beta = 2 - m_0^\beta$. The random vector M can thus take four possible values, whose probabilities are parameterized by the matrix $(p_{i,j}) = (\mathbb{P}\{M = (m_i^\alpha, m_j^\beta)\})_{0 \leq i,j \leq 1}$. The conditions $\mathbb{P}(M^\alpha = m_0^\alpha) = 1/2$ and $\mathbb{P}(M^\beta = m_0^\beta) = 1/2$ impose that

$$\begin{bmatrix} p_{00} & p_{01} \\ p_{10} & p_{11} \end{bmatrix} = \begin{bmatrix} \frac{1 + \rho_m^*}{4} & \frac{1 - \rho_m^*}{4} \\ \frac{1 - \rho_m^*}{4} & \frac{1 + \rho_m^*}{4} \end{bmatrix}$$

for some $\rho_m^* \in [-1, 1]$. We easily check that ρ_m^* is the correlation between components M^α and M^β under the distribution M . The empirical section focuses on the specification $\rho_m^* = 1$, which in unreported work is never rejected on the currency data. Because distributions with $\rho_m^* < 1$ may be useful for series with weaker correlation in volatility, we report in the next subsection all theoretical results for arbitrary values of ρ_m^* .

3.2. Properties

The dynamics of $M_{k,t}$ are determined by the switching vector $1_{k,t}$ and the bivariate distribution M . We show in the Appendix A that $M_{k,t}$ has a unique ergodic distribution $\bar{\Pi}_k$. Let $\bar{\Pi}_k^{HH}$, $\bar{\Pi}_k^{HL}$, $\bar{\Pi}_k^{LH}$ and $\bar{\Pi}_k^{LL}$, respectively, denote the probabilities of states (m_0^α, m_0^β) , (m_0^α, m_1^β) , (m_1^α, m_0^β) , and (m_1^α, m_1^β) . The symmetry of the construction implies that $\bar{\Pi}_k^{HH} = \bar{\Pi}_k^{LL}$ and $\bar{\Pi}_k^{HL} = \bar{\Pi}_k^{LH}$. When $\rho_m^* > 0$, the multipliers are more likely to be either both high or both low: $\bar{\Pi}_k^{HH} = \bar{\Pi}_k^{LL} > 1/4 > \bar{\Pi}_k^{HL} = \bar{\Pi}_k^{LH}$.

Since different components are statistically independent, the ergodic distribution of the volatility state $M_t = (M_{1,t}; \dots; M_{\bar{k},t})$ is the product measure $\bar{\Pi} = \bar{\Pi}_1 \otimes \dots \otimes \bar{\Pi}_{\bar{k}}$. Under this distribution, the cross-currency correlation of volatility components is given by

$$\text{Corr}(M_{kt}^\alpha, M_{kt}^\beta) = \rho_m^* \frac{(1 - \lambda)\gamma_k + \lambda}{2 - [(1 - \lambda)\gamma_k + \lambda]} \leq \rho_m^*.$$

This coefficient is large when the transition probability γ_k or the correlation of switching events λ are high, i.e. when arrivals tend to happen at the same time. The upper bound ρ_m^* is reached when either $\lambda = 1$ or $\gamma_k \rightarrow 1$.

The return series have unconditional means equal to zero: $\mathbb{E}x_t = 0$. By the Cauchy–Schwarz inequality, their correlation satisfies

$$\text{Corr}(x_t^\alpha; x_t^\beta) = \rho_\varepsilon \prod_{k=1}^{\bar{k}} \mathbb{E}[(M_{k,t}^\alpha M_{k,t}^\beta)^{1/2}] \leq \rho_\varepsilon. \tag{3.3}$$

The upper bound ρ_ε is attained when the multipliers of both series are perfectly correlated. On the other hand when $\rho_\varepsilon < 1$, uncorrelated changes in volatility represent additional sources of noise that reduce the correlation of asset returns.

The econometrician does not observe the volatility state, but only the set of past returns $X_t \equiv \{x_s\}_{s=1}^t$. Returns are unpredictable with this information set: $\mathbb{E}_{t-1}x_t = 0$, and the model is thus consistent with some standard forms of market efficiency.¹⁰ Comovement is quantified by the conditional covariance $\text{Cov}_t(x_{t+n}^\alpha; x_{t+n}^\beta) = \rho_\varepsilon \sigma_\alpha \sigma_\beta \prod_{k=1}^{\bar{k}} \mathbb{E}_t[(M_{k,t+n}^\alpha M_{k,t+n}^\beta)^{1/2}]$, and the conditional correlation

$$\text{Corr}_t(x_{t+n}^\alpha; x_{t+n}^\beta) = \rho_\varepsilon \prod_{k=1}^{\bar{k}} \frac{\mathbb{E}_t[(M_{k,t+n}^\alpha M_{k,t+n}^\beta)^{1/2}]}{[(\mathbb{E}_t M_{k,t+n}^\alpha)(\mathbb{E}_t M_{k,t+n}^\beta)]^{1/2}} \leq \rho_\varepsilon. \tag{3.4}$$

These quantities fluctuate through time with the multipliers. Thus while the construction assumes constant correlation coefficients ρ_ε and ρ_m , the conditional correlation of returns is time-varying. We easily check that it is large when the volatility components of the currencies are high.

Comovement in volatility can be similarly investigated. We observe that when $\rho_m > 0$ and \bar{k} is large, the conditional correlation of absolute returns satisfies

$$\text{Corr}_t(|x_{t+n}^\alpha|; |x_{t+n}^\beta|) \sim C_\varepsilon \prod_{k=1}^{\bar{k}} \frac{\mathbb{E}_t[(M_{k,t+n}^\alpha M_{k,t+n}^\beta)^{1/2}]}{[(\mathbb{E}_t M_{k,t+n}^\alpha)(\mathbb{E}_t M_{k,t+n}^\beta)]^{1/2}}, \tag{3.5}$$

where $C_\varepsilon = \mathbb{E}(|\varepsilon_1^\alpha \varepsilon_1^\beta|)$. Consistent with previous intuition, correlation between absolute returns is high in periods of high volatility.

4. Inference

We now develop inference methods for bivariate MSM. Bayesian updating and the likelihood function are available in closed-form, and ML estimation is therefore practical with a moderate number of volatility components. Alternative computational methods are designed for high-dimensional state spaces.

4.1. Closed-form likelihood

Since each frequency vector M_{kt} is drawn from a bivariate binomial, the volatility state $M_t = (M_{1,t}; M_{2,t}; \dots; M_{\bar{k},t})$ takes $d = 4^{\bar{k}}$ possible values $m^1, \dots, m^d \in \mathbb{R}_+^{\bar{k}}$. The

¹⁰See Campbell et al. (Chapter 2, 1997) for a discussion.

dynamics of M_t are characterized by the $d \times d$ transition matrix $A = (a_{ij})_{1 \leq i,j \leq d}$ with components $a_{ij} = \mathbb{P}(M_{t+1} = m^j | M_t = m^i)$.

The econometrician does not observe the volatility state, but only the set of past returns $X_t \equiv \{x_s\}_{s=1}^t$. The corresponding probabilities

$$\Pi_t^j = \mathbb{P}(M_t = m^j | X_t)$$

are computed recursively by Bayesian updating. Let $\Pi_t = (\Pi_t^1, \dots, \Pi_t^d) \in \mathbb{R}_+^d$. In the next period, state M_{t+1} is drawn and the econometrician observes the return vector x_{t+1} . Conditional on the new volatility state, the joint return x_{t+1} has bivariate Gaussian density $f_{x_{t+1}}(x_{t+1} | M_{t+1} = m^j)$. By Bayes' rule, the updated probability is

$$\Pi_{t+1} = \frac{f(x_{t+1}) * \Pi_t A}{[f(x_{t+1}) * \Pi_t A] \mathbf{1}'}, \tag{4.1}$$

where $\mathbf{1} = (1, \dots, 1) \in \mathbb{R}^d$ and $f(x)$ is the vector of conditional densities $(f_{x_{t+1}}(x_{t+1} | M_{t+1} = m^j))_j$. The belief Π_{t+1} is thus a function of the contemporaneous observation x_{t+1} and the prior probability Π_t . In empirical applications, the initial vector Π_0 is chosen equal to the ergodic distribution $\bar{\Pi}$ of the Markov chain.

We easily check that the log-likelihood function has the closed-form expression:

$$\ln L(x_1, \dots, x_T; \psi) = \sum_{t=1}^T \ln [f(x_t) \cdot (\Pi_{t-1} A)].$$

For fixed \bar{k} , the ML estimator is consistent and asymptotically efficient as $T \rightarrow \infty$. Analytical multistep forecasting can proceed using updated beliefs and the transition matrix A as in Calvet and Fisher (2001).

4.2. Particle filter

The transition matrix A contains $4^{\bar{k}} \times 4^{\bar{k}}$ elements, increasing exponentially in the number of frequencies. When $\bar{k} = 8$, the transition matrix thus has cardinality $2^{32} \approx 4 \times 10^9$, and is computationally expensive to use. Following the recent literature on Markov chains,¹¹ a simulation-based inference methodology is proposed. Specifically, we introduce a particle filter, a recursive algorithm that generates independent draws $M_t^{(1)}, \dots, M_t^{(B)}$ from the conditional distribution Π_t .

We begin at $t = 0$ by drawing $M_0^{(1)}, \dots, M_0^{(B)}$ from the ergodic distribution $\bar{\Pi}$. For any $t \geq 0$, assume that $\{M_t^{(b)}\}_{b=1}^B$ have been independently sampled from Π_t . Given a new return x_{t+1} , an approximation to Bayes' rule generates draws $\{M_{t+1}^{(b)}\}_{b=1}^B$ from the new belief Π_{t+1} . More specifically, we rewrite the updating formula (4.1) as

$$\Pi_{t+1}^j \propto f_{x_{t+1}}(x_{t+1} | M_{t+1} = m^j) \sum_{i=1}^d \mathbb{P}(M_{t+1} = m^j | M_t = m^i) \Pi_t^i.$$

¹¹See for instance Chib et al. (2002), Jacquier et al. (1994), and Pitt and Shephard (1999).

The vectors $M_t^{(1)}, \dots, M_t^{(B)}$ are independent draws from Π_t . This suggests the Monte Carlo approximation:

$$\Pi_{t+1}^j \propto f_{x_{t+1}}(x_{t+1}|M_{t+1} = m^j) \frac{1}{B} \sum_{b=1}^B \mathbb{P}(M_{t+1} = m^j|M_t = M_t^{(b)}).$$

As shown in the Appendix A, we complete the approximation by simulating each $M_t^{(b)}$ one-step forward and reweighting using an importance sampler:

1. Simulate the Markov chain one-step ahead to obtain $\hat{M}_{t+1}^{(1)}$ given $M_t^{(1)}$. Repeat B times to generate B draws $\hat{M}_{t+1}^{(1)}, \dots, \hat{M}_{t+1}^{(B)}$. This preliminary step only uses information available at date t , and must therefore be adjusted to account for the new return.
2. Draw a random number q from 1 to B with probability

$$\mathbb{P}(q = b) \equiv \frac{f_{x_{t+1}}(x_{t+1}|M_{t+1} = \hat{M}_{t+1}^{(b)})}{\sum_{b'=1}^B f_{x_{t+1}}(x_{t+1}|M_{t+1} = \hat{M}_{t+1}^{(b')})}.$$

The vector $M_{t+1}^{(1)} = \hat{M}_{t+1}^{(q)}$ is a draw from Π_{t+1} . Repeat B times to obtain B draws $M_{t+1}^{(1)}, \dots, M_{t+1}^{(B)}$.

This recursive procedure provides a discrete approximation to Bayesian updating, which is computationally convenient in large state spaces.

4.3. Simulated likelihood

We can use the particle filter to compute the likelihood function. Each one-step ahead density satisfies $f(x_t|X_{t-1}) = \sum_{i=1}^d f(x_t|M_t = m^i)\mathbb{P}(M_t = m^i|X_{t-1})$. Given simulated draws $\hat{M}_t^{(b)}$ from $M_t|X_{t-1}$, the Monte Carlo estimate of the conditional density is defined as

$$\hat{f}(x_t|X_{t-1}) \equiv \frac{1}{B} \sum_{b=1}^B f(x_t|M_t = \hat{M}_t^{(b)})$$

and the log-likelihood is approximated by $\sum_{t=1}^T \ln \hat{f}(x_t|X_{t-1})$. We can use these calculations to carry out simulated likelihood estimation. In practice, an arbitrarily close approximation can be achieved by increasing B . Larger state spaces require more draws to achieve the same degree of precision.

Table 5 presents a Monte Carlo assessment of this method. We focus on the univariate specification with $\bar{k} = 8$ components. Using the particle filter, we generate 500 approximations of the log-likelihood of the univariate DM series at the optimized ML estimates from Table 1. Each calculation uses independent sets of Monte Carlo draws. We then compare the mean, standard deviation, and quantiles of the estimates with the exact value obtained in Table 1 by analytical Bayesian updating. All particle filter evaluations use $B = 10,000$ random draws. The particle filter estimate of the log-likelihood has a relatively small standard deviation and the

Table 5
Evaluation of particle filter

	ln L	$E_t \sum_{j=1}^n x_{t+j}^2$				
		$n = 1$	5	10	20	50
True value	-6885.9	0.432	2.194	4.442	8.991	22.66
Simulation average	-6887.3	0.431	2.187	4.426	8.953	22.54
Standard deviation	1.8	0.012	0.064	0.142	0.325	0.98
1% quantile	-6892.1	0.405	2.036	4.076	8.103	19.79
25% quantile	-6888.4	0.423	2.147	4.338	8.747	21.97
50% quantile	-6887.3	0.431	2.191	4.435	8.957	22.66
75% quantile	-6886.2	0.439	2.231	4.525	9.179	23.23
99% quantile	-6883.3	0.458	2.330	4.739	9.633	24.42

This table compares values generated by the particle filter with their true values generated by exact Bayesian updating. $\ln L$ is the value of the log-likelihood function for the Deutsche Mark series with $\bar{k} = 8$ evaluated at the maximum likelihood estimates in Table 1. The forecasted variance of the series is denoted $E_t \sum_{j=1}^n x_{t+j}^2$. For each quantity, the table provides the true value along with the average, standard deviation, and quantiles over 500 particle filter approximations using independent sets of random draws. Each approximation uses $B = 10,000$ random draws.

average across simulations, -6887.3 , is close to the true value of -6885.9 . The quantiles are tightly clustered as well. The table also shows particle filter estimates of the forecast variance, which are accurate and approximately unbiased. These results confirm that the particle filter produces reasonable estimates of the likelihood and moments of the series for problems of reasonable size.

The particle filter extends the range of computationally feasible multifractal specifications. In previous work with univariate processes, Calvet and Fisher (2004) report an approximate computational upper bound of 10 binomial components, or 2^{10} states. While this gives good results in the univariate case, multivariate models require a correspondingly larger number of state variables. We will thus show in the empirical section that the particle filter produces good results in a bivariate model with $\bar{k} = 8$ components, or 2^{16} states. The particle filter also permits implementation of specifications where the state vector M is drawn from a continuous distribution. Since earlier research (Calvet and Fisher, 2002) suggest that exchange rates are best fit by lognormal multipliers, it will be interesting in future work to revisit the lognormal specification and compare its performance to the binomial.

4.4. Two-step estimation

Two-step estimation offers additional computational benefits, permitting the econometrician to decompose inference into a sequence of lower-dimensional optimization problems. In the bivariate multifractal, each series $c \in \{\alpha, \beta\}$ follows a univariate MSM with parameters m_0^c, σ_c, b and $\gamma_{\bar{k}}^c$. This implies that we can estimate six of the eight parameters using the likelihood and smaller state space of the univariate model. Additionally, univariate estimation gives good precision in finite

samples (Calvet and Fisher, 2004). This motivates us to develop the two-step method described below. The Appendix A shows that this procedure is a special case of GMM, implying consistency and asymptotic normality of the estimator.

In the first stage, we obtain the parameters $(m_0^\alpha, m_0^\beta, \sigma_\alpha, \sigma_\beta, b, \gamma_{\bar{k}})$ by optimizing the sum of the two univariate log-likelihoods, as in (2.2). Intuitively, this gives consistent estimates for all parameters since the gradient of this sum with respect to the true parameters is zero. Because this objective function coincides with the likelihood of the combined univariate, the first step has already been completed in Section 2.

The second stage gives estimates for the remaining two parameters, $(\rho_\varepsilon, \lambda)$, which are unique to the bivariate model. When the state space is not too large, $(\bar{k} \leq 5)$, computation of the analytical bivariate likelihood is practical. We therefore maximize the exact bivariate MSM probability density conditional on the first-stage estimates. For higher-dimensional specifications, $(\bar{k} = 6, 7, 8)$, computation of the analytical bivariate likelihood is difficult. We therefore use the particle filter to optimize the simulated likelihood as described in Section 4.3.¹² In this paper, the two-step procedure aids empirical implementation of bivariate specifications with state spaces as large as 2^{16} .

5. Empirical results

5.1. Bivariate MSM estimates

We report in Table 6 empirical ML estimates for $\bar{k} \leq 5$ and exchange rate pairs (DM, JA), (DM, UK) and (JA, UK). The estimates are obtained by maximizing the full likelihood of bivariate MSM, as described in Section 4.1. As in univariate MSM, \hat{m}_0 declines with \bar{k} , while the standard deviations $\hat{\sigma}_\alpha$ and $\hat{\sigma}_\beta$ tend to vary with \bar{k} with no apparent trend. The correlation between Gaussian innovations $\hat{\rho}_\varepsilon$ is positive and roughly constant across \bar{k} . The arrival correlation $\hat{\lambda}$ is also large and approximately invariant to the number of volatility components. The two parameters specific to the bivariate specification, $\hat{\rho}_\varepsilon$ and $\hat{\lambda}$, thus seem precisely estimated. Finally, the estimated $\hat{\lambda}$ is highest when $\hat{\rho}_\varepsilon$ is highest, and lowest when $\hat{\rho}_\varepsilon$ is lowest. We infer that correlation in volatility is higher for currencies with more correlated returns, which is economically intuitive.

The likelihood functions sharply increase with the number of frequencies. For instance with (DM, JA), the log-likelihood increases by more than 800 when \bar{k} goes from 1 to 5. Since the models are non-nested and specified by the same number of parameters, this is a very substantial increase of fit in-sample. We also compare the goodness of fit to the independent case in Table 3, and find that for (DM, JA) with

¹²One could of course more generally match to any relevant moments in the second stage. Our view is that simulated likelihood is an excellent choice for intermediate problems because it potentially entails a small loss in efficiency. For very large problems, including many assets as discussed in Section 6, it would be natural to consider matching moments such as (3.3) and (3.5). This could potentially further reduce computational requirements.

Table 6
 Bivariate MSM: full maximum likelihood estimates

	$\bar{k} = 1$	2	3	4	5
<i>DM and JA</i>					
\hat{m}_0^{DM}	1.637 (0.011)	1.589 (0.013)	1.543 (0.013)	1.484 (0.013)	1.447 (0.011)
\hat{m}_0^{JA}	1.718 (0.011)	1.701 (0.009)	1.667 (0.010)	1.621 (0.010)	1.573 (0.010)
$\hat{\sigma}_{\text{DM}}$	0.679 (0.009)	0.621 (0.011)	0.575 (0.014)	0.559 (0.017)	0.524 (0.015)
$\hat{\sigma}_{\text{JA}}$	0.683 (0.011)	0.649 (0.014)	0.577 (0.017)	0.573 (0.018)	0.509 (0.024)
$\hat{\gamma}_{\bar{k}}$	0.122 (0.013)	0.217 (0.022)	0.732 (0.066)	0.828 (0.049)	0.905 (0.038)
\hat{b}	—	16.23 (3.09)	23.71 (4.54)	13.60 (1.48)	8.70 (0.83)
$\hat{\rho}_\varepsilon$	0.580 (0.008)	0.589 (0.009)	0.576 (0.010)	0.580 (0.009)	0.580 (0.009)
$\hat{\lambda}$	0.647 (0.041)	0.641 (0.039)	0.589 (0.056)	0.634 (0.048)	0.637 (0.049)
$\ln L$	-12519.99	-12001.70	-11797.05	-11688.44	-11655.80
<i>DM and UK</i>					
\hat{m}_0^{DM}	1.651 (0.012)	1.570 (0.010)	1.522 (0.012)	1.492 (0.012)	1.484 (0.011)
\hat{m}_0^{UK}	1.731 (0.010)	1.656 (0.010)	1.624 (0.009)	1.588 (0.017)	1.564 (0.010)
$\hat{\sigma}_{\text{DM}}$	0.681 (0.009)	0.706 (0.014)	0.626 (0.011)	0.560 (0.031)	0.498 (0.012)
$\hat{\sigma}_{\text{UK}}$	0.629 (0.009)	0.658 (0.015)	0.573 (0.012)	0.506 (0.042)	0.458 (0.015)
$\hat{\gamma}_{\bar{k}}$	0.227 (0.021)	0.422 (0.052)	0.746 (0.057)	0.791 (0.067)	0.864 (0.040)
\hat{b}	—	13.29 (2.28)	15.24 (2.26)	11.71 (1.68)	10.83 (1.35)
$\hat{\rho}_\varepsilon$	0.707 (0.007)	0.714 (0.007)	0.707 (0.007)	0.708 (0.007)	0.710 (0.007)
$\hat{\lambda}$	0.837 (0.023)	0.852 (0.023)	0.833 (0.026)	0.822 (0.027)	0.827 (0.025)
$\ln L$	-10894.41	-10513.18	-10335.82	-10270.90	-10240.51
<i>JA and UK</i>					
\hat{m}_0^{JA}	1.764 (0.014)	1.718 (0.008)	1.693 (0.009)	1.629 (0.010)	1.608 (0.010)
\hat{m}_0^{UK}	1.729 (0.005)	1.661 (0.012)	1.633 (0.012)	1.595 (0.011)	1.571 (0.010)
$\hat{\sigma}_{\text{JA}}$	0.655 (0.008)	0.619 (0.014)	0.531 (0.015)	0.489 (0.014)	0.709 (0.021)
$\hat{\sigma}_{\text{UK}}$	0.603 (0.006)	0.578 (0.012)	0.514 (0.018)	0.474 (0.011)	0.385 (0.009)

Table 6 (continued)

	$\bar{k} = 1$	2	3	4	5
$\hat{\gamma}_{\bar{k}}$	0.219 (0.011)	0.304 (0.027)	0.449 (0.054)	0.748 (0.046)	0.791 (0.043)
\hat{b}	—	21.50 (4.32)	15.08 (2.08)	13.21 (1.43)	11.91 (1.40)
$\hat{\rho}_\varepsilon$	0.447 (0.007)	0.453 (0.004)	0.449 (0.011)	0.438 (0.012)	0.440 (0.011)
$\hat{\lambda}$	0.499 (0.048)	0.565 (0.047)	0.560 (0.054)	0.544 (0.056)	0.535 (0.059)
$\ln L$	-12247.45	-11647.36	-11404.09	-11266.91	-11211.52

This table shows maximum likelihood estimation results for bivariate MSM. Columns correspond to the number \bar{k} of volatility components. Asymptotic standard errors are in parenthesis.

$\bar{k} = 5$, the gain in likelihood is over 1300 points. Results are similar for other currencies, demonstrating that the bivariate model improves over independent univariate models.

In Table 7, we re-estimate bivariate MSM with the two-step procedure of Section 4.4. The second stage uses the analytical bivariate likelihood for $\bar{k} \leq 5$, and the particle filter for $\bar{k} \in \{6, 7, 8\}$. We observe that with $\bar{k} \leq 5$, the two-step parameter estimates are comparable to the full MLE results, which suggests that the two-step procedure works well. For $\bar{k} \in \{6, 7, 8\}$, the results appear consistent with the univariate MLE of Table 1 as well as the estimates of the lower dimensional bivariate models. The particle filter is thus effective in extending the range of tractable models.

We compare bivariate MSM with the constant correlation GARCH (CC-GARCH) of Bollerslev (1990), which is a standard benchmark in the multivariate volatility literature. Returns are specified as

$$x_t^\alpha = \sqrt{h_t^\alpha} \varepsilon_t^\alpha, \quad x_t^\beta = \sqrt{h_t^\beta} \varepsilon_t^\beta,$$

where ε_t^α and ε_t^β are two standard normals with correlation ρ_ε . The conditional variances h_t^α and h_t^β satisfy the recursions $h_{t+1}^c = \omega_c + a_c(\varepsilon_t^c)^2 + b_c h_t^c$ for each $c \in \{\alpha, \beta\}$. CC-GARCH is thus specified by 7 parameters as compared to 8 with multivariate MSM.

We report in Table 8 an in-sample comparison of CC-GARCH against bivariate MSM with $\bar{k} = 5$ components. It is immediately clear that MSM gives much higher likelihoods although it has only one additional parameter. For all three pairs of exchange rates, the difference in log-likelihood exceeds 1000 points. The same results hold whether comparing full MLE results from the two models, or the likelihoods obtained under two-step estimation. To account for the difference in the number of parameters, we compute the BIC statistic for each model. We then test the significance of the difference using the original method suggested by Vuong (1989), and the HAC-adjusted version developed in Calvet and Fisher (2004). In all cases,

Table 7
Bivariate MSM: two-step estimates

	$k = 1$	2	3	4	5	6	7	8
<i>DM and JA</i>								
\hat{m}_0^1	1.643 (0.020)	1.618 (0.019)	1.515 (0.022)	1.474 (0.023)	1.445 (0.022)	1.405 (0.022)	1.397 (0.022)	1.367 (0.022)
\hat{m}_0^2	1.775 (0.013)	1.757 (0.012)	1.687 (0.016)	1.638 (0.017)	1.578 (0.020)	1.565 (0.018)	1.522 (0.019)	1.488 (0.021)
$\hat{\sigma}_1$	0.669 (0.014)	0.577 (0.011)	0.597 (0.019)	0.569 (0.020)	0.504 (0.021)	0.565 (0.018)	0.449 (0.027)	0.472 (0.035)
$\hat{\sigma}_2$	0.613 (0.010)	0.544 (0.010)	0.565 (0.018)	0.487 (0.016)	0.476 (0.021)	0.632 (0.017)	0.384 (0.022)	0.532 (0.041)
$\hat{\gamma}_{\bar{k}}$	0.129 (0.014)	0.257 (0.024)	0.301 (0.037)	0.756 (0.072)	0.844 (0.075)	0.872 (0.081)	0.959 (0.047)	0.982 (0.027)
\hat{b}	—	69.57 (21.80)	11.97 (2.20)	13.21 (2.11)	9.14 (1.32)	7.16 (1.29)	6.16 (0.86)	4.93 (0.56)
$\hat{\rho}_\varepsilon$	0.566 (0.013)	0.570 (0.014)	0.581 (0.016)	0.574 (0.017)	0.578 (0.017)	0.581 (0.049)	0.581 (0.009)	0.618 (0.010)
$\hat{\lambda}$	0.587 (0.065)	0.544 (0.067)	0.646 (0.064)	0.585 (0.082)	0.624 (0.080)	0.633 (0.032)	0.659 (0.038)	0.633 (0.023)
<i>DM and UK</i>								
\hat{m}_0^1	1.626 (0.020)	1.565 (0.021)	1.519 (0.022)	1.473 (0.023)	1.452 (0.022)	1.406 (0.022)	1.401 (0.023)	1.370 (0.022)
\hat{m}_0^2	1.697 (0.016)	1.657 (0.017)	1.641 (0.019)	1.602 (0.021)	1.573 (0.022)	1.521 (0.022)	1.492 (0.022)	1.454 (0.022)
$\hat{\sigma}_1$	0.671 (0.013)	0.645 (0.017)	0.599 (0.018)	0.568 (0.021)	0.493 (0.018)	0.563 (0.018)	0.471 (0.028)	0.474 (0.033)
$\hat{\sigma}_2$	0.605 (0.011)	0.588 (0.015)	0.515 (0.015)	0.468 (0.016)	0.422 (0.017)	0.457 (0.020)	0.391 (0.019)	0.385 (0.026)
$\hat{\gamma}_{\bar{k}}$	0.090 (0.011)	0.151 (0.019)	0.388 (0.052)	0.683 (0.082)	0.672 (0.087)	0.798 (0.093)	0.844 (0.092)	0.969 (0.043)
\hat{b}	—	12.33 (3.38)	15.25 (3.02)	11.97 (2.07)	10.09 (1.68)	6.97 (1.48)	6.23 (0.88)	5.02 (0.65)
$\hat{\rho}_\varepsilon$	0.697 (0.010)	0.703 (0.011)	0.704 (0.012)	0.709 (0.013)	0.711 (0.012)	0.700 (0.011)	0.689 (0.012)	0.704 (0.010)
$\hat{\lambda}$	0.814 (0.041)	0.818 (0.037)	0.826 (0.042)	0.802 (0.048)	0.820 (0.045)	0.790 (0.050)	0.844 (0.022)	0.800 (0.027)
<i>JA and UK</i>								
\hat{m}_0^1	1.776 (0.062)	1.762 (0.031)	1.688 (0.032)	1.645 (0.025)	1.631 (0.024)	1.568 (0.024)	1.558 (0.021)	1.507 (0.023)
\hat{m}_0^2	1.728 (0.016)	1.680 (0.016)	1.640 (0.019)	1.607 (0.020)	1.575 (0.021)	1.525 (0.022)	1.499 (0.021)	1.458 (0.021)
$\hat{\sigma}_1$	0.624 (0.011)	0.546 (0.011)	0.569 (0.017)	0.471 (0.014)	0.702 (0.028)	0.631 (0.030)	0.514 (0.025)	0.509 (0.031)
$\hat{\sigma}_2$	0.606 (0.012)	0.561 (0.012)	0.522 (0.015)	0.508 (0.015)	0.432 (0.016)	0.457 (0.021)	0.384 (0.018)	0.379 (0.023)

Table 7 (continued)

	$k = 1$	2	3	4	5	6	7	8
$\hat{\gamma}_{\bar{k}}$	0.161 (0.015)	0.303 (0.027)	0.275 (0.030)	0.635 (0.062)	0.697 (0.068)	0.847 (0.064)	0.864 (0.067)	0.970 (0.034)
\hat{b}	—	43.46 (10.62)	12.73 (2.27)	14.55 (2.14)	13.60 (2.08)	8.22 (1.04)	7.24 (0.91)	5.60 (0.70)
$\hat{\rho}_\varepsilon$	0.439 (0.017)	0.439 (0.018)	0.448 (0.019)	0.439 (0.021)	0.439 (0.021)	0.436 (0.010)	0.414 (0.018)	0.436 (0.017)
$\hat{\lambda}$	0.494 (0.068)	0.519 (0.071)	0.570 (0.063)	0.549 (0.076)	0.524 (0.080)	0.575 (0.049)	0.625 (0.027)	0.561 (0.015)

This table shows two-step estimates for bivariate MSM. Columns correspond to the number \bar{k} of volatility components. First stage estimates are obtained by optimizing the combined univariate likelihood as in Table 3A. As described in the Appendix A, this provides consistent estimates for the parameters $(m_0^\alpha, m_0^\beta, \sigma_\alpha, \sigma_\beta, b, \gamma_{\bar{k}})$. For $\bar{k} \leq 5$, the second stage optimizes the analytically calculated bivariate MSM likelihood conditional on the first stage estimates. For $\bar{k} = 6, 7, 8$, the optimization of the likelihood is numerically implemented using the particle filter approximation. Standard errors in parentheses are calculated by recasting the optimization in a GMM context, as described in the Appendix A, and are HAC adjusted using Newey and West (1987).

the p -value from the test that CC-GARCH has a superior BIC statistic to multivariate MSM is substantially less than 1%. In-sample evidence thus strongly favors multivariate MSM.

5.2. Integral transforms

We now perform out-of-sample diagnostics with probability integral transforms, as in Diebold et al. (1998) and Elerian et al. (2001). In all remaining empirical work, we consider the MSM specification with $\bar{k} = 5$ components. We first estimate MSM and CC-GARCH on the 1973–1989 subsample. The out-of-sample evaluations are based on the daily 3473 observations from 1990 to 2003. Let $y_{t,n} \equiv \sum_{i=1}^n x_{t+i}$ denote the forward-looking n -period return at time t . MSM and CC-GARCH each generate a conditional forecast distribution

$$F_{t,n}(y) \equiv \mathbb{P}(y_{t,n} \leq y | x_1, \dots, x_t).$$

Under correct specification, the random variables $U_{t,n} = F_{t,n}(y_{t,n})$ are uniformly distributed on $[0, 1]$; they are also independent if $n = 1$.

In Fig. 1, we compare histograms of selected integral transforms $\{U_{t,n}\}$ for the two models. Histograms are shown for $n = 1$ and $n = 5$ days using the following portfolios: DM, JA, an equal-weighted position in the two currencies, and a hedge portfolio with weights $(1, -1)$.¹³ We see that MSM provides approximately uniform

¹³The random variables U_t are constructed as follows. In every period, we use the particle filter to draw B values $y_{t,n}^{(1)}, \dots, y_{t,n}^{(B)}$ from the conditional distribution of $y_{t,n}$ given x_1, \dots, x_t . We then approximate $F_{t,n}(y)$ by the empirical c.d.f. $\hat{F}_{t,n}(y) = 1/B \sum_{b=1}^B 1\{y_{t,n}^{(b)} \leq y\}$. Sensitivity tests indicate that $B = 10,000$ draws are more than sufficient to provide a good approximation.

Table 8
In-sample model comparison

	No. of Parameters	ln <i>L</i>	BIC	BIC <i>p</i> -value vs. Multifractal	
				Vuong (1989)	HAC Adj
<i>A. Full MLE estimates</i>					
DM and JA					
Bivariate MSM	8	−11655.80	3.0626		
CC GARCH	7	−12825.63	3.3679	<0.001	<0.001
DM and UK					
Bivariate MSM	8	−10240.51	2.6919		
CC GARCH	7	−11331.00	2.9764	<0.001	<0.001
JA and UK					
Bivariate MSM	8	−11211.52	2.9462		
CC GARCH	7	−12550.49	3.2958	<0.001	<0.001
<i>B. Two-step estimates</i>					
DM and JA					
Bivariate MSM	8	−11658.89	3.0634		
CC GARCH	7	−12830.98	3.3693	<0.001	<0.001
DM and UK					
Bivariate MSM	8	−10262.05	2.6975		
CC GARCH	7	−11434.17	3.0034	<0.001	<0.001
JA and UK					
Bivariate MSM	8	−11233.59	2.9521		
CC GARCH	7	−12559.72	3.2982	<0.001	<0.001

This table summarizes information about in-sample goodness of fit. The Bayesian information criterion is given by $BIC = T^{-1}(-2 \ln L + NP \ln T)$. The last two columns give *p*-values from a test that the corresponding model dominates the multifractal model by the BIC criterion. The first value uses the Vuong (1989) methodology, and the second value adjusts the test for heteroskedasticity and autocorrelation as described in Calvet and Fisher (2004). A low *p*-value indicates that the CC GARCH model would be rejected in favor of the multifractal model. Panel A presents the results when both models have been estimated by Full MLE. Panel B presents results where both models are estimated by two-step procedures.

histograms. In contrast, CC-GARCH generates tent-shaped plots with a large concentration of values around 0 and 1. These feature are symptomatic of tails that are too thin in the estimated CC-GARCH process. Similar results are obtained with other currencies.

The CVM criterion confirms these graphical results. Let T^* denote the number of out-of-sample periods, and \hat{F}_U the empirical distribution of the transforms $U_{t,1}$. As

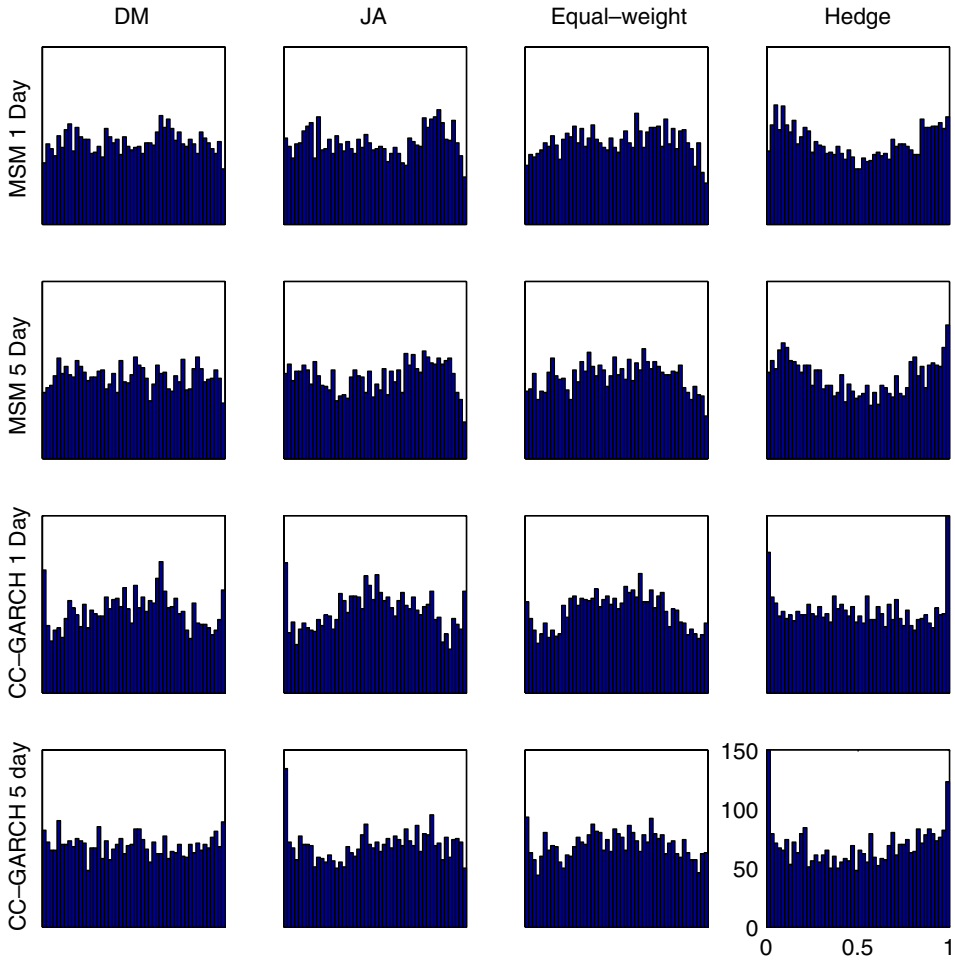


Fig. 1. Probability integral transforms. These figures show histograms of the probability integral transforms $\{U_{t,n}\}$ for horizons (in rows) of $n = 1$ and $n = 5$ days and portfolios (in columns) of DM, JA, an equal-weighted portfolio of the two currencies, and a hedge portfolio with weights $(1, -1)$. The models considered are bivariate MSM with $\bar{k} = 5$ components and CC-GARCH. Under correct specification, the integral transforms are uniformly distributed.

$T^* \rightarrow \infty$, the CVM criterion $T^* \int_0^1 [y - \hat{F}_U(y)]^2 dy$ weakly converges to a weighted series of independent χ^2 random variables:

$$T^* \int_0^1 [y - \hat{F}_U(y)]^2 dx \Rightarrow \sum_{j=1}^{\infty} \left(\frac{z_j}{j\pi} \right)^2,$$

where the $\{z_j\}$ are IID $\mathcal{N}(0, 1)$.¹⁴ We report in Table 9 the CVM statistics for all currencies and portfolios. At the 1% level, we reject MSM in only 2 out of 12 cases,

¹⁴See Shorack and Wellner (1986) for further details.

Table 9
Goodness of fit 1-day forecasts

	Bivariate MSM			CC GARCH		
	DM, JA	DM, UK	JA, UK	DM, JA	DM, UK	JA, UK
Currency α	0.22	0.16	0.54	1.07	1.07	1.70
Currency β	0.47	0.51	0.29	1.70	2.57	2.56
Equal-weight	0.67	0.05	0.23	2.48	1.06	2.51
Hedge	2.27	0.59	1.06	0.56	5.93	0.52

This table shows the CVM distance between a uniform distribution and the empirical distribution of the probability integral transform of the corresponding model forecast. The bivariate MSM specification uses $\bar{k} = 5$ components. Currency α and β refer to the first currency and second currency in each pair. Equal-weight is an equal weighted portfolio, and Hedge is a zero investment portfolio of $\alpha - \beta$. Under correct specification, the reported statistics are greater than 0.73 in about 1% of samples. A high value of the statistic thus indicates rejection of the corresponding model. Rejections at the 1% level are indicated by bold face.

while CC-GARCH is rejected 10 out 12 cases.¹⁵ The CVM statistics thus confirm that CC-GARCH provides inaccurate conditional density forecasts, while MSM is broadly consistent with exchange rate data.

5.3. Value-at-risk

The tail properties of financial series are of direct interest for risk management and financial regulation. VaR is a particularly widespread method that summarizes the expected maximum loss over a target horizon within a given confidence interval. Given a confidence level p , we define the VaR of a portfolio to be the $1 - p$ th quantile of the conditional return distribution:

$$VaR_{t,n}(p) \equiv F_{t,n}^{-1}(1 - p).$$

Thus with probability p we expect to lose no more than $VaR_{t,n}(p)$ over the next n days.

The accuracy of a VaR model is most easily verified by recording the failure rate, i.e. the number of times VaR is exceeded in a given sample (e.g., Kupiec, 1995; Jorion, 1997).¹⁶ Table 9 reports the failure rates of MSM and CC-GARCH for portfolios held for $n = 1$ and 5 days and confidence levels of 1%, 5% and 10%. As described in the Appendix A, we forecast for each bivariate process the VaR of individual currencies, equal-weighted portfolios, and hedge portfolios.

The results in Table 10 lead to two conclusions: MSM is more conservative and more accurate than CC-GARCH. MSM is more conservative because it tends to fail less than CC-GARCH. For example, when the 1-day predicted failure rate is 1%,

¹⁵We do not adjust the critical values for estimation error. Earlier work (e.g., Thompson, 2000) suggests that such adjustments would only have small effects.

¹⁶The failure rate is thus the proportion of out-of-sample days in which $x_{t+1} + \dots + x_{t+n} < VaR_{t,n}(p)$.

Table 10
Failure rates value at risk forecasts

	Bivariate MSM			CC GARCH		
	1%	5%	10%	1%	5%	10%
<i>A. 1-day horizon</i>						
DM and JA						
Currency α	0.69	4.35	9.10	1.81	5.13	9.01
Currency β	0.95	4.81	9.56	2.30	5.38	9.10
Equal-weight	0.86	3.92	8.32	1.30	4.66	8.21
Hedge	0.69	5.64	12.21	2.25	6.68	11.81
DM and UK						
Currency α	0.92	4.92	10.14	1.81	5.13	9.01
Currency β	0.72	5.27	10.68	1.44	4.61	8.29
Equal-weight	1.07	4.69	10.28	1.87	5.18	8.98
Hedge	0.55	4.72	9.13	0.92	4.00	7.00
JA and UK						
Currency α	1.01	5.04	9.88	2.30	5.38	9.10
Currency β	0.60	4.41	9.70	1.44	4.61	8.29
Equal-weight	0.84	4.55	8.78	1.64	4.69	8.03
Hedge	1.15	5.64	11.37	2.25	6.25	10.34
<i>B. 5-day horizon</i>						
DM and JA						
Currency α	0.78	4.21	9.57	1.61	5.48	10.72
Currency β	1.07	5.30	10.55	2.31	7.06	11.62
Equal-weight	0.72	4.44	8.50	1.64	5.25	9.14
Hedge	0.92	5.42	12.16	3.29	8.39	13.46
DM and UK						
Currency α	0.95	5.13	10.35	1.64	5.51	10.72
Currency β	0.75	5.28	11.01	0.98	4.79	9.77
Equal-weight	0.84	5.07	10.93	1.27	5.94	10.61
Hedge	0.69	4.01	8.76	0.86	3.86	8.48
JA and UK						
Currency α	1.21	5.53	10.75	2.36	6.86	11.70
Currency β	0.46	4.67	9.02	1.53	4.93	9.57
Equal-weight	0.84	4.12	9.02	1.53	4.93	9.57
Hedge	1.73	6.40	11.24	1.76	6.34	11.53

This table displays the frequency of returns that exceed the VaR forecasted by the model. The bivariate MSM specification uses $\bar{k} = 5$ components. For quantile $p\%$ the number reported is the frequency of portfolio returns below quantile p predicted by the model. If the VaR forecast is correct, the observed failure rate should be close to the prediction. Boldface numbers are statistically different from p at the 1% level. Panel A shows results for a one day horizon, while Panel B shows a 5-day horizon. Currency α and β refer to the first currency and second currency in each pair. Equal-weight is an equal weighted portfolio, and Hedge is a zero investment portfolio of $\alpha - \beta$. Standard errors in Panel A are computed by $p(1-p)/3473$, where 3473 is the number of out-of-sample observations. Standard errors in Panel B are computed using Newey and West (1987).

actual portfolio losses exceed the MSM VaR forecast more than 1% of the time for 3 out of 12 portfolios. Actual losses exceed the 1% CC-GARCH quantile more than 1% of the time in 11 out of 12 portfolios. Of course, an excessively conservative model does not necessarily lead to superior risk management. Statistical tests suggest that MSM is not overly conservative. For each portfolio and VaR quantile we test the null hypothesis that the empirical failure rate is equal to the expected failure rate. For the MSM model, the failure rates are statistically different from the 1% prediction for only 1 out of 12 portfolios. The CC-GARCH failure rates are statistically different from 1% in 11 out of 12 portfolios. MSM thus provides more accurate quantile forecasts than CC-GARCH.

5.4. Specification tests

Comparisons in- and out-of-sample have shown that bivariate MSM performs well relatively to CC-GARCH. It is now natural to investigate whether, in an absolute sense, the restrictions imposed by our model are supported by the data. We weaken one assumption at a time, and assess improvement in fit by LR tests. When a restriction applies equally to the univariate and bivariate models, we choose to test on the univariate series. This allows us to distinguish between misspecifications originating in univariate MSM and those unique to the bivariate approach.

Heterogeneity in volatility persistence is made parsimonious by the frequency parameterization (3.2). We focus on univariate models with $\bar{k} = 8$ components. For each currency, we consider the *restricted* univariate MLE estimates in Table 1 and denote by L_r the corresponding likelihood. In contrast, we call *unrestricted* model $k \in \{1, \dots, \bar{k}\}$ the extension in which frequency parameter γ_k is free and all other frequencies satisfy (3.2). We estimate the k th unrestricted model and denote by $L_u(k)$ the corresponding likelihood. Under the restricted model, $2[\ln L_u(k) - \ln L_r]$ converges to $\chi^2(1)$ as $T \rightarrow \infty$. This methodology generates eight LR statistics for each of the three currencies. For space constraints, we report in the text only the salient features of the analysis. For DM, none of the tests provides evidence against the MSM frequency restrictions at the 1% level. One statistic ($k = 6$) is significant for JA, and two tests ($k = 6, 7$) are significant for UK. Evidence against the frequency specification is thus limited to three of the twenty-four tests.

We similarly assess on univariate series whether volatility components have identical distributions across frequencies. Unrestricted specification k permits component $M_{k,t}$ to have its own distribution parameter $m_0(k)$. Results are mixed. For DM, only two of the eight tests ($k = 1, 6$) are significant at the 1% level. For JA, the first five tests suggest a value of $m_0(k)$ larger than for the other components. Similarly for the UK series, LR tests of the first 4 components suggest stronger shocks at low than at high-frequency variation, while higher frequencies ($k = 6, 7$) suggest less variability. Overall, the DM data seems to match the MSM model exceptionally well, while JA and UK appear to prefer stronger low-frequency variation. These results are consistent with Calvet and Fisher (2002), who show that

DM best matches the moment-scaling restrictions implied by the multifractal model. We also note that the current analysis only considers binomial MSM. Multivariate distributions M such as the lognormal may better accommodate the strong low-frequency variations in UK and JA over the last three decades.

We finally test the restrictions imposed by bivariate MSM on volatility comovement. For each currency pair, the restricted model is given by the full MLE estimates with $\bar{k} = 5$ in Table 6. Unrestricted model k permits that component k may have its own unique arrival correlation λ_k . We report no rejections at the 1% level for JA–UK, one significant test ($k = 5$) for the DM–UK pair, and two significant statistics ($k = 2, 5$) for the DM–JA pair. Overall, MSM incorporates empirically reasonable restrictions that permit parsimonious specification of bivariate multifrequency volatility.

6. Extension to many assets

6.1. Multivariate MSM

Bivariate MSM can be readily extended to economies with an arbitrary number N of financial prices, as is now shown. The construction assumes in every period a volatility component $M_{k,t}^n \in \mathbb{R}_+$ for each frequency $k \in \{1, \dots, \bar{k}\}$ and asset $n \in \{1, \dots, N\}$. As in the bivariate case, components $M_{k,t}^n$ and $M_{k',t}^{n'}$ can be correlated across assets ($n \neq n'$), but are statistically independent across frequencies ($k \neq k'$). Specifically, the volatility dynamics are determined by a fixed multivariate distribution M on \mathbb{R}_+^N , and an arrival vector $1_{k,t} \in \{0, 1\}^N$ for each frequency. The k th component of every asset switches with unconditional probability γ_k ($\mathbb{E}1_{k,t} = \gamma_k \mathbf{1}$), and arrivals across assets are characterized by the correlation matrix:

$$\text{Corr}(1_{k,t}) = (\lambda_{n,n'})_{1 \leq n, n' \leq N}.$$

The state vector is defined recursively. At time t , we draw the independent arrival vector $(1_{k,t})_{k=1, \dots, \bar{k}}$, and sample the new components $M_{k,t}$ from the corresponding marginal distribution of M .

The volatility state is fully specified by $N \times \bar{k}$ matrix $M_t = (M_{k,t}^n)_{n,k}$. The econometrician has beliefs Π_t over the state space that can be updated using Bayes' rule. When the distribution of M is discrete, the likelihood function is available in closed-form. For large state spaces, estimation can be carried out using a particle filter. We refer the reader to the Appendix A for further details.

While natural, this approach requires the specification and estimation of the multivariate distribution M and the arrival correlation matrix $(\lambda_{n,n'})_{1 \leq n, n' \leq N}$. In a general formulation, the number of parameter therefore grows at least as fast as a quadratic function of N . Like other specifications such as multivariate GARCH, the model is computationally expensive for a large number of assets. The next subsection develops an overlapping class of models, which is based on the same principles as multivariate MSM and yet remains tractable with many assets.

6.2. Factor MSM

We now propose a factor model for stochastic volatility, which generates a parsimonious multifrequency specification. The construction is based on L volatility factors $F_t^\ell = (F_{k,t}^\ell)_{1 \leq k \leq \bar{k}} \in \mathbb{R}_+^{\bar{k}}$ that can jointly affect all assets. For instance, the vector F_t^ℓ may contain the frequency-specific components determining the volatility of a global risk factor or a specific industry.¹⁷ Each vector F_t^ℓ contains \bar{k} frequency-specific components and follows a specific univariate MSM process with parameters $(b, \gamma_{\bar{k}}, m_0^\ell)$. The volatility of each asset n is also affected by an idiosyncratic shock E_t^n , which is specified by parameters $(b, \gamma_{\bar{k}}, m_0^{L+n})$. Draws of the factors $F_{k,t}^\ell$ and idiosyncratic shocks $E_{k,t}^n$ are independent, but the timing of arrivals may be correlated. Factors and idiosyncratic components thus follow univariate MSM with identical frequencies.

For every asset n and frequency k , the volatility component $M_{k,t}^n$ is the weighted product of the factors and idiosyncratic shock of same frequency:

$$M_{k,t}^n = C_n (F_{k,t}^1)^{w_1^n} \dots (F_{k,t}^L)^{w_L^n} (E_{k,t}^n)^{w_{L+1}^n}.$$

The weights are non-negative and add up to one. The constant C_n is chosen to guarantee that $\mathbb{E}M_{k,t}^n = 1$, and is thus not a free parameter.¹⁸ In logarithms, we obtain the familiar additive formulation:

$$\ln M_{k,t}^n = \ln C_n + \sum_{\ell=1}^L w_\ell^n \ln F_{k,t}^\ell + w_{L+1}^n \ln E_{k,t}^n.$$

Returns are then defined as previously as $x_t = (M_{1,t} * \dots * M_{\bar{k},t})^{1/2} * \varepsilon_t$, where ε_t is a centered multivariate Gaussian noise: $\varepsilon_t \sim N(0, \Sigma)$. We show in the Appendix A that as with multivariate GARCH, the estimation of Σ can be carried out directly from sample autocorrelations of the series.

Two special cases of this setup are of particular interest. First, when arrivals for all factors and idiosyncratic components are simultaneous, factor MSM is a special case of the multivariate MSM in the previous subsection. New draws of $M_{k,t}^n$ are then independent of all past multipliers, and the factor model generates univariate series that are consistent with univariate MSM. Further, when the distribution of factors and idiosyncratic shocks is lognormal, the resulting multipliers $M_{k,t}^n$ are lognormal as well. This is convenient as we know that lognormal multipliers fit well the moment-scaling properties of financial series, including exchange rates (Calvet and Fisher, 2002). Stochastic volatility is now fully specified by: (1) the frequency parameters b and $\gamma_{\bar{k}}$; (2) the distribution parameters of factors and idiosyncratic shocks $(m_0^1, \dots, m_0^{L+N})$; and (3) the factor loadings $w^n = (w_1^n, \dots, w_L^n)$ of each asset. The model is thus defined by $N(L + 1) + L + 2$ volatility parameters.

¹⁷See Diebold and Rudebusch (1996) for an early discussion of the macroeconomic motivation underlying Markov-switching factor models.

¹⁸We thus have $C_n = 1 / [\mathbb{E}[(F_{k,t}^1)^{w_1^n}] \dots \mathbb{E}[(F_{k,t}^L)^{w_L^n}] \mathbb{E}[(E_{k,t}^n)^{w_{L+1}^n}]]$. This computation is particularly easy when the marginal distribution of the shocks are multinomial or lognormal.

The second interesting special case is when arrivals of factors and idiosyncratic shocks are independent. It is easy to verify that this specification has the same number of parameters as when arrivals are simultaneous. Further, this choice permits that at time t some but not all factors and idiosyncratic components may change. The univariate volatility components $M_{k,t}^n$ then takes a new value without requiring a completely independent draw from M . Thus, the implied univariate volatility dynamics are smoother than standard MSM, but can generate the same thick tails and long-memory volatility persistence. These specifications are thus both practical to implement and deserving of further empirical investigation.

7. Conclusion

This paper uses the Markov-switching multifractal (MSM) of Calvet and Fisher (2001, 2004) to implement a univariate frequency decomposition of volatility in several exchange rate series. We find that the estimated components are generally difficult to relate to standard macroeconomic variables. Low-frequency volatility components from all currencies covary positively with oil and gold prices, suggesting that these commodities may act as proxies for global economic risk. Relative to previous measures of volatility, the frequency decomposition increases the strength and cross-sectional robustness of our results. At the same time, our analysis encourages the econometrician to be cautious since the source of covariation is limited to low frequencies.

We identify strong patterns in volatility comovement between currencies. Across exchange-rate pairs, volatility components tend to have high correlation when their durations are similar and low correlations otherwise. This motivates the development of bivariate MSM, a multifrequency model of comovement in stochastic volatility and covariation in financial prices. The model permits a parsimonious specification of bivariate shocks with heterogeneous durations, capturing the economic intuition that shared fundamentals may have different innovation frequencies. Bayesian updating and the likelihood function are always available in closed-form, but are practical only when the state space is of moderate size. We develop a particle filter suitable for larger state spaces, and demonstrate its good performance in inference and forecasting. We estimate bivariate MSM on three exchange rate pairs, and show that it performs well in- and out-of-sample relative to a standard benchmark, CC-GARCH. LR tests also confirm some of the principal restrictions of the model. We conduct inference and forecasting for good performing pure regime-switching models with 2^{16} states and only eight parameters.

The methods developed in this paper open a number of new directions for future research. First, MSM offers an economically appealing and computationally tractable alternative to previous multivariate GARCH and stochastic volatility models. Comparisons of the approaches in different applications can and should be developed. Additionally, the particle filter methodology opens new frontiers for conducting estimation and inference in MSM processes with very large state spaces. The particle filter also permits examination of processes where the volatility

component distribution takes values on a continuous support. Earlier work suggests that lognormal distributions might be particularly appealing. This specification becomes computationally accessible with the particle filter, and can be compared to the binomial specification used in this paper and earlier research. Finally, we propose in this paper a multifrequency factor structure appropriate for multivariate settings with potentially large numbers of assets, which appears promising for future empirical research.

Acknowledgements

We received helpful comments from Frank Diebold (the editor), two anonymous referees, John Campbell, Ken Rogoff, Chris Sims, Robert Stambaugh, Jim Stock, Mungo Wilson and seminar participants at Harvard University and the PIER-IGIER Conference on Econometric Methods in Macroeconomics and Finance. This paper is a revised version of an earlier manuscript, “Comovement and Predictor Variables for Multifractal Volatility” (Calvet et al., 2003). We are very appreciative of financial support provided for this project by the Social Sciences and Humanities Research Council of Canada under Grant 410-2002-0641.

Appendix A

A.1. Distribution of the arrival vector

The probability of a simultaneous switch is $\mathbb{P}(1_{k,t}^\alpha = 1 | 1_{k,t}^\beta = 1) = \mathbb{P}(1_{k,t}^\alpha = 1 | 1_{k,t}^\beta = 1 | 1_{k,t}^\alpha = 1)$. Overall, the vector $1_{k,t}$ has joint distribution

	Arrival on β	No arrival on β
Arrival on α	$\gamma_k[(1 - \lambda)\gamma_k + \lambda]$	$\gamma_k(1 - \lambda)(1 - \gamma_k)$
No arrival on α	$\gamma_k(1 - \gamma_k)(1 - \lambda)$	$(1 - \gamma_k)[1 - \gamma_k(1 - \lambda)]$

A.2. Ergodic distribution of volatility components

The bivariate process $(M_{k,t}^\alpha, M_{k,t}^\beta)$ can take values $s^1 = (H, H)$, $s^2 = (H, L)$, $s^3 = (L, H)$ and $s^4 = (L, L)$. The transition matrix is $T = (t_{ij})$, where $t_{ij} = \mathbb{P}(s_{t+1} = s^j | s_t = s^i)$. It satisfies

$$T = \begin{bmatrix} p_k & 1 - \frac{\gamma_k}{2} - p_k & 1 - \frac{\gamma_k}{2} - p_k & \gamma_k - 1 + p_k \\ 1 - \frac{\gamma_k}{2} - q_k & q_k & \gamma_k - 1 + q_k & 1 - \frac{\gamma_k}{2} - q_k \\ 1 - \frac{\gamma_k}{2} - q_k & \gamma_k - 1 + q_k & q_k & 1 - \frac{\gamma_k}{2} - q_k \\ \gamma_k - 1 + p_k & 1 - \frac{\gamma_k}{2} - p_k & 1 - \frac{\gamma_k}{2} - p_k & p_k \end{bmatrix},$$

where

$$p_k = 1 - \gamma_k + \gamma_k[(1 - \lambda)\gamma_k + \lambda] \frac{1 + \rho_m^*}{4},$$

$$q_k = 1 - \gamma_k + \gamma_k[(1 - \lambda)\gamma_k + \lambda] \frac{1 - \rho_m^*}{4}.$$

Simple manipulation implies that the characteristic polynomial of T is

$$P_T(x) = (1 - x)(1 - \gamma_k - x)^2[2(p_k + q_k + \gamma_k) - 3 - x].$$

We easily check that $|2(p_k + q_k + \gamma_k) - 3| < 1$ and infer that T has a unique ergodic distribution $\bar{\Pi}_k = (\bar{\Pi}_k^{HH}, \bar{\Pi}_k^{HL}, \bar{\Pi}_k^{LH}, \bar{\Pi}_k^{LL})$. The symmetry of the transition matrix implies that $\bar{\Pi}_k^{HH} = \bar{\Pi}_k^{LL}$ and $\bar{\Pi}_k^{HL} = \bar{\Pi}_k^{LH}$. We easily check that $\bar{\Pi}_k^{HH} = \frac{1}{4} \frac{2 - 2q_k - \gamma_k}{2 - (p_k + q_k) - \gamma_k}$ or equivalently

$$\bar{\Pi}_k^{HH} = \frac{1}{4} \frac{1 - (1 - \rho_m^*)[(1 - \lambda)\gamma_k + \lambda]/2}{1 - [(1 - \lambda)\gamma_k + \lambda]/2},$$

and finally note that $\bar{\Pi}_k^{HL} = 1/2 - \bar{\Pi}_k^{HH}$.

A.3. Particle filter

As discussed in the main text, the vectors $\hat{M}_{t+1}^{(1)}, \dots, \hat{M}_{t+1}^{(B)}$ are independent draws from the probability distribution $h(m) \equiv \mathbb{P}(M_{t+1} = m | X_t)$. Consider a continuous function Y defined on \mathbb{R}_+^k and taking values on the real line. The conditional expectation $\mathbb{E}[Y(M_{t+1}) | X_{t+1}] = \sum_{j=1}^d \mathbb{P}(M_{t+1} = m^j | X_{t+1}) Y(m^j)$ is conveniently re-written as

$$\mathbb{E}[Y(M_{t+1}) | X_{t+1}] = \sum_{j=1}^d h(m^j) \frac{\mathbb{P}(M_{t+1} = m^j | X_{t+1})}{h(m^j)} Y(m^j).$$

The Monte Carlo approximation to this integral is

$$\mathbb{E}[Y(M_{t+1}) | X_{t+1}] \approx \frac{1}{B} \sum_{b=1}^B \frac{\mathbb{P}(M_{t+1} = \hat{M}_{t+1}^{(b)} | X_{t+1})}{h(\hat{M}_{t+1}^{(b)})} Y(\hat{M}_{t+1}^{(b)}).$$

Bayes' rule implies

$$\frac{\mathbb{P}(M_{t+1} = \hat{M}_{t+1}^{(b)} | X_{t+1})}{B h(\hat{M}_{t+1}^{(b)})} = \frac{f_{x_{t+1}}(x_{t+1} | M_{t+1} = \hat{M}_{t+1}^{(b)})}{B f_{x_{t+1}}(x_{t+1} | X_t)}.$$

Since $f_{x_{t+1}}(x_{t+1} | X_t) \approx \frac{1}{B} \sum_{b'=1}^B f_{x_{t+1}}(x_{t+1} | \hat{M}_{t+1}^{(b')})$, we infer that

$$\frac{\mathbb{P}(M_{t+1} = \hat{M}_{t+1}^{(b)} | X_{t+1})}{B h(\hat{M}_{t+1}^{(b)})} \approx \frac{f_{x_{t+1}}(x_{t+1} | M_{t+1} = \hat{M}_{t+1}^{(b)})}{\sum_{b'=1}^B f_{x_{t+1}}(x_{t+1} | M_{t+1} = \hat{M}_{t+1}^{(b')})}.$$

The right-hand side defines a probability μ_b for every $b \in \{1, \dots, B\}$. We infer that the random variable $Y(M_{t+1})$ has conditional expectation $\mathbb{E}[Y(M_{t+1}) | X_{t+1}] \approx$

$\sum_{b=1}^B \mu_b Y(\hat{M}_{t+1}^{(b)})$. Since this result is valid for any function Y , we conclude that Π_{t+1} can be approximated with a discrete distribution taking on the value $\hat{M}_{t+1}^{(b)}$ with probability μ_b .

A.4. Two-step estimation

We partition the parameter vector into $\psi \equiv (\psi'_1, \psi'_2)'$, with $\psi_1 = (\sigma_\alpha, \sigma_\beta, m_0^\alpha, m_0^\beta, b, \gamma_{\bar{k}})'$ and $\psi_2 = (\rho_\varepsilon, \lambda)'$. In the first step, we compute the vector $\hat{\psi}_1$ that maximizes the combined univariate likelihood $L(x_t^\alpha; m_0^\alpha, \sigma_\alpha, b, \gamma_{\bar{k}}) + L(x_t^\beta; m_0^\beta, \sigma_\beta, b, \gamma_{\bar{k}})$. Note that $\hat{\psi}_1$ is a GMM estimator based on the moment conditions $\partial L(x_t^\alpha)/\partial \psi_1 + \partial L(x_t^\beta)/\partial \psi_1$. Under correct specification the expectation of each derivative is zero, which implies consistency and asymptotic normality of $\hat{\psi}_1$. In the second step, we estimate ψ_2 by maximizing the simulated bivariate likelihood $L(x_t^\alpha, x_t^\beta; \hat{\psi}_1, \psi_2)$ given the first stage estimate $\hat{\psi}_1$. The simulated likelihood is computed using the particle filter with $B = 10,000$ draws.

Standard errors for the two-step estimates are obtained by restating the algorithm as a GMM estimator based on the moment conditions $T^{-1} \sum_{t=1}^T g_t(\hat{\psi}) = 0$, where $g_t(\psi)$ is the column vector with components $\partial[\ln f(x_t^\alpha | X_{t-1}^\alpha) + \ln f(x_t^\beta | X_{t-1}^\beta)]/\partial \psi_1$ and $\partial \ln f(x_t^\alpha, x_t^\beta | X_{t-1}^\alpha, X_{t-1}^\beta)/\partial \psi_2$. Standard GMM arguments imply asymptotic normality

$$\sqrt{T}(\hat{\psi} - \psi_0) \xrightarrow{d} \mathcal{N}[0, H^{-1}V(H')^{-1}]$$

with $H = -\mathbb{E} \partial g_t(\psi_0)/\partial \psi'$ and $V = \text{Var}[T^{-1/2} \sum g_t(\psi_0)]$. To estimate V , we approximate g_t by taking finite difference derivatives of the objective function. Then we estimate V using the formula of Newey and West (1987) with 10 lags. When calculating finite difference derivatives using the particle filter, we use 15,000 simulations. We estimate H by calculating the sample variance of the first derivatives:

$$\hat{H} = \begin{pmatrix} \hat{H}_{1,1} + \hat{H}_{1,2} & 0 \\ \hat{H}_{2,1} & \hat{H}_{2,2} \end{pmatrix},$$

where $\hat{H}_{1,1}$ and $\hat{H}_{1,2}$ are the 6×6 matrices

$$\hat{H}_{1,1} = T^{-1} \sum \frac{\partial \ln f(x_t^\alpha | X_{t-1}^\alpha)}{\partial \psi_1} \frac{\partial \ln f(x_t^\alpha | X_{t-1}^\alpha)}{\partial \psi_1'} \approx -\mathbb{E} \left[\frac{\partial^2 \ln f(x_t^\alpha | X_{t-1}^\alpha)}{\partial \psi_1 \partial \psi_1'} \right],$$

$$\hat{H}_{1,2} = T^{-1} \sum \frac{\partial \ln f(x_t^\beta | X_{t-1}^\beta)}{\partial \psi_1} \frac{\partial \ln f(x_t^\beta | X_{t-1}^\beta)}{\partial \psi_1'} \approx -\mathbb{E} \left[\frac{\partial^2 \ln f(x_t^\beta | X_{t-1}^\beta)}{\partial \psi_1 \partial \psi_1'} \right].$$

Similarly, $(\hat{H}_{21}, \hat{H}_{22})$ are the bottom two rows of the 8×8 matrix

$$T^{-1} \sum \frac{\partial \ln f(x_t^\alpha, x_t^\beta | X_{t-1}^\alpha, X_{t-1}^\beta)}{\partial \psi} \frac{\partial \ln f(x_t^\alpha, x_t^\beta | X_{t-1}^\alpha, X_{t-1}^\beta)}{\partial \psi'}.$$

The matrix \hat{H} is consistent since its elements are second derivatives of the univariate or bivariate likelihoods.

A.5. Value-at-risk forecasts

We use the particle filter to calculate the VaR implied by MSM. The algorithm in Section 3.4 is used to generate volatility draws $M_t^{(1)}, \dots, M_t^{(B)}$ from the distribution Π_t . For each draw $M_t^{(b)}$, we simulate the bivariate series forward n days to obtain B draws from the cumulative return on the portfolio. We then estimate $VaR_{t,n}(p)$ as the $1 - p$ th empirical quantile.

For 1 day forecasts CC-GARCH provides a closed form expression for VaR, namely $VaR_{t,1}(p) = -Q_{1-p}\sigma_{t|t-1}$, where Q_{1-p} is the $(1 - p)$ th quantile of a standard normal variable and $\sigma_{t|t-1}$ is the standard deviation implied by CC-GARCH. The 5-day CC-GARCH forecasts are calculated by simulation. In all cases we use $B = 10,000$ simulated draws.

A.6. Inference in the multivariate model

For either multivariate MSM or factor MSM, we seek to estimate the covariance matrix Σ and the vector of volatility parameters ψ . One possibility is to choose a tight specification for Σ and use the particle filter to optimize the simulated likelihood over all parameters. For example, our bivariate estimates show that currency pairs with strong volatility comovement also have high correlation in innovations. This suggests using the same factor structure that controls volatility to parsimoniously specify Σ .

In the general case, estimation can be conducted in two steps: (1) Estimate the covariance matrix of the Gaussian noises; (2) Use the particle filter to estimate the volatility parameters ψ by simulated ML. Step (2) is straightforward, and step (1) can be conducted as follows. For any two assets n and p , we know that

$$\mathbb{E}[x_t^{(n)} x_t^{(p)}] = \Gamma_{n,p} \mathbb{E}[\varepsilon_t^{(n)} \varepsilon_t^{(p)}] \quad \text{and} \quad \mathbb{E}[|x_t^{(n)} x_t^{(p)}|] = \Gamma_{n,p} \mathbb{E}[|\varepsilon_t^{(n)} \varepsilon_t^{(p)}|],$$

where $\Gamma_{n,p} = \prod_{k=1}^k \mathbb{E}\{[M_{k,t}^n M_{k,t}^p]^2\}^{1/2}$. We infer

$$\frac{\sum_t x_t^{(n)} x_t^{(p)}}{\sum_t |x_t^{(n)} x_t^{(p)}|} \xrightarrow{\text{a.s.}} \frac{\mathbb{E}[\varepsilon_t^{(n)} \varepsilon_t^{(p)}]}{\mathbb{E}[|\varepsilon_t^{(n)} \varepsilon_t^{(p)}|]} = \varphi(\rho_{n,p}),$$

where $\rho_{n,p} = \text{Corr}[\varepsilon_t^{(n)}; \varepsilon_t^{(p)}]$ and $\varphi(\rho) \equiv \frac{\pi}{2} \frac{\rho}{\sqrt{1 - \rho^2 + \rho \arcsin \rho}}$. The function φ is strictly increasing and maps $[-1, 1]$ onto $[-1, 1]$. A consistent estimator of the correlation coefficient is therefore

$$\hat{\rho}_{n,p} = \varphi^{-1} \left(\frac{\sum_t x_t^{(n)} x_t^{(p)}}{\sum_t |x_t^{(n)} x_t^{(p)}|} \right).$$

The variance of the Gaussians is consistently estimated by $\hat{\sigma}_n^2 = \frac{1}{T} \sum_t [x_t^{(n)}]^2$. The covariance matrix defined by $\hat{\sigma}_n^2$ and $\hat{\rho}_{n,p}$ may not be positive-definite. We thus apply the methodology of Ledoit et al. (2003) to obtain a positive semi-definite matrix $\hat{\Sigma}$.

References

- Andersen, T., Bollerslev, T., 1998. Deutsche-mark dollar volatility: intraday activity patterns, macroeconomic announcements, and longer-run dependencies. *Journal of Finance* 53, 219–265.
- Andersen, T., Bollerslev, T., Lange, S., 1999. Forecasting financial market volatility: sample frequency vis-à-vis forecast horizon. *Journal of Empirical Finance* 6, 457–477.
- Andersen, T., Benzoni, L., Lund, J., 2002. An empirical investigation of continuous-time models for equity returns. *Journal of Finance* 57, 1239–1284.
- Bollerslev, T., 1990. Modelling the coherence in short-run nominal exchange rates: a multivariate generalized ARCH approach. *Review of Economics and Statistics* 72, 498–505.
- Bollerslev, T., Engle, R., Wooldridge, J., 1988. A capital asset pricing model with time varying covariances. *Journal of Political Economy* 96, 116–131.
- Calvet, L., Fisher, A., 2001. Forecasting multifractal volatility. *Journal of Econometrics* 105, 27–58.
- Calvet, L., Fisher, A., 2002. Multifractality in asset returns: theory and evidence. *Review of Economics and Statistics* 84, 381–406.
- Calvet, L., Fisher, A., 2004. How to forecast long-run volatility: regime-switching and the estimation of multifractal processes. *Journal of Financial Econometrics* 2, 49–83.
- Calvet, L., Fisher, A., Thompson, S., 2003. Comovement and predictor variables for multifractal volatility. Harvard and UBC Working Paper.
- Campbell, J., Lo, A., MacKinlay, C., 1997. *The Econometrics of Financial Markets*. Princeton University Press, Princeton, NJ.
- Chib, S., Nardari, F., Shephard, N., 2002. Markov chain Monte Carlo methods for stochastic volatility models. *Journal of Econometrics* 108, 281–316.
- Clarida, R., Sarno, L., Taylor, M., Valente, G., 2003. The out-of-sample success of term structure models as exchange rate predictors: a step beyond. *Journal of International Economics* 60, 61–83.
- Diebold, F., Inoue, A., 2001. Long memory and regime switching. *Journal of Econometrics* 105, 27–58.
- Diebold, F., Nerlove, M., 1989. The dynamics of exchange rate volatility: a multivariate latent factor ARCH model. *Journal of Applied Econometrics* 4, 1–21.
- Diebold, F., Rudebusch, G., 1996. Measuring business cycles: a modern perspective. *Review of Economics and Statistics* 78, 67–77.
- Diebold, F., Gunther, T., Tay, A., 1998. Evaluating density forecasts with applications to financial risk management. *International Economic Review* 39, 863–883.
- Ding, Z., Granger, C., Engle, R., 1993. A long memory property of stock market returns and a new model. *Journal of Empirical Finance* 1, 83–106.
- Edwards, S., Susmel, S., 2003. Interest-rate volatility in emerging markets. *Review of Economics and Statistics* 85, 328–348.
- Elerian, O., Chib, S., Shephard, N., 2001. Likelihood inference for discretely observed nonlinear diffusions. *Econometrica* 69, 959–993.
- Engle, R., 1987. Multivariate GARCH with factor structures—cointegration in variance. Working Paper, Department of Economics, University of California, San Diego.
- Engle, R., 2002. Dynamic conditional correlation: a simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business and Economic Statistics* 20, 339–350.
- Engle, R., Kroner, K., 1995. Multivariate simultaneous generalized ARCH. *Econometric Theory* 11, 122–150.
- Engle, R., Mezrich, J., 1996. GARCH for groups. *Risk* 9, 36–40.
- Engle, R., Ito, T., Lin, W.-L., 1990a. Meteor showers or heat waves? Heteroskedastic intra-daily volatility in the foreign exchange market. *Econometrica* 58, 525–542.
- Engle, R., Ng, V., Rothschild, M., 1990b. Asset pricing with a factor ARCH covariance structure: empirical estimates for treasury bills. *Journal of Econometrics* 45, 213–238.
- Hamilton, J.D., 1989. A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica* 57, 357–384.
- Hamilton, J.D., 2003. What is an oil shock? *Journal of Econometrics* 113, 363–398.

- Harvey, A., Ruiz, E., Shephard, N., 1994. Multivariate stochastic variance models. *Review of Economic Studies* 61, 247–264.
- Hidalgo, J., Robinson, P.M., 1996. Testing for structural change in a long-memory environment. *Journal of Econometrics* 70, 159–174.
- Jacquier, E., Polson, N., Rossi, P., 1994. Bayesian analysis of stochastic volatility models. *Journal of Business and Economic Statistics* 12, 371–417.
- Johannes, M., Polson, N., Stroud, J., 2002. Sequential optimal portfolio performance: market and volatility timing. Working Paper, University of Chicago Graduate School of Business.
- Jorion, P., 1997. *Value At Risk: The New Benchmark for Controlling Market Risk*. McGraw Hill, New York.
- Kim, C.J., 1993. Dynamic linear models with Markov-switching. *Journal of Econometrics* 60, 1–22.
- Kraft, D., Engle, R., 1982. Autoregressive conditional heteroskedasticity in multiple time series. Unpublished manuscript, Department of Economics, University of California, San Diego.
- Kupiec, P., 1995. Techniques for verifying the accuracy of risk measurement models. *Journal of Derivatives* 2, 73–84.
- Ledoit, O., Santa-Clara, P., Wolf, M., 2003. Flexible multivariate GARCH modeling with an application to international stock markets. *Review of Economics and Statistics* 85, 735–747.
- Lyons, R., 1995. Tests of microstructural hypotheses in the foreign exchange market. *Journal of Financial Economics* 39, 321–351.
- Lyons, R., 2001. *The Microstructure Approach to Exchange Rates*. MIT Press, Cambridge, MA.
- Meese, R., Rogoff, K., 1983. Empirical exchange rate models of the seventies: do they fit out of sample? *Journal of International Economics* 14, 3–24.
- Newey, W., West, K., 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica* 55, 703–708.
- Pitt, M., Shephard, N., 1999. Filtering via simulation: auxiliary particle filter. *Journal of the American Statistical Association* 94, 590–599.
- Rogoff, K., 1999. Perspectives on exchange rate volatility. In: Feldstein, M. (Ed.), *International Capital Flows*. University of Chicago Press, Chicago, pp. 441–453.
- Sarno, L., Taylor, M., 2002. *The Economics of Exchange Rates*. Cambridge University Press, Cambridge.
- Schwert, G.W., 1989. Why does stock market volatility change over time? *Journal of Finance* 44, 1115–1153.
- Shorack, G., Wellner, J., 1986. *Empirical Processes with Applications to Statistics*. Wiley, New York.
- Stock, J., Watson, M., 2003. Has the business cycle changed? Evidence and explanations. Working Paper, Harvard University.
- Thompson, S.B., 2000. Specification tests for continuous time models. University of California at Berkeley, Ph.D. Dissertation (Chapter 3).
- Vuong, Q., 1989. Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica* 57, 307–333.