

A Bayesian analysis of market information linkages among NAFTA countries using a multivariate stochastic volatility model

Petra Fleischer · Ross Maller · Gernot Müller

Published online: 28 May 2009
© Springer Science + Business Media, LLC 2009

Abstract NAFTA has arguably been the most important and elaborate free-trade agreement in history, providing a blueprint for potential new agreements. So far, the evidence is mixed as to whether NAFTA has been successful in terms of its economic impact. We fit a multivariate stochastic volatility model that directly measures financial information linkages across the three participating countries in a trivariate setting. The model detects significant changes in information linkages across the countries from the pre- to post-NAFTA period with a high degree of reliability. This has implications not only for measuring these linkages but also for hedging and portfolio diversification

This research was partially supported by ARC grant DP0664603.

P. Fleischer · R. Maller
School of Finance and Applied Statistics,
College of Business and Economics,
The Australian National University, Canberra,
Australian Capital Territory 0200, Australia
e-mail: Petra.Fleischer@anu.edu.au

R. Maller
Centre for Mathematics and its Applications,
The Australian National University, Canberra,
Australian Capital Territory 0200, Australia
e-mail: Ross.Maller@anu.edu.au

G. Müller (✉)
Centre of Mathematics, Technische Universität München,
85747 Garching, Germany
e-mail: mueller@ma.tum.de

policies. An MCMC procedure is used to fit the model, and the accuracy and robustness of the method is confirmed by simulations.

Keywords North American Free Trade Agreement (NAFTA) · Information and Volatility Linkages · Volatility Correlations · Markov Chain Monte Carlo · Equity Market Returns

JEL Classifications C11 · C32 · F13

1 Introduction

The North American Free Trade Agreement (NAFTA) came into effect from 1 January 1994. It is an agreement to largely reduce or completely eliminate tariffs and other trade barriers on products and services between the three participating countries: the US, Canada and Mexico. To date, NAFTA's influence has been examined extensively for international trade and globalisation aspects. We do not attempt to review this literature here,¹ but, in brief, some evidence suggests that NAFTA has been generally successful in promoting trade among the US, Canada and Mexico, cf. Gould (1996, 1998). This would lead us to expect an increase in equity market linkages between the signature parties following the introduction of NAFTA. Stiglitz (2004), on the other hand, reports that NAFTA appeared to have failed to reduce income disparities between the US and Mexico. In fact, rather than reducing those disparities, he reports an increase of 10.6% in them over the previous ten years. Thus, evidence as to whether or not the NAFTA countries have indeed become more closely linked is mixed, so far.

Some detailed empirical studies, likewise, have produced an ambiguous picture. A cointegration analysis by Ewing et al. (1999) found no evidence of integration in the index returns of the US, Canada and Mexico over the period 1987–1997. Darrat and Zhong (2005) extended the cointegration analysis to the longer period June 1, 1989 through April 10, 2002, augmented by a profile analysis to investigate the existence of an equilibrium, and found fairly robust evidence indicating intensified equity market linkage since the NAFTA accord. The longer period and larger data set for analysis no doubt increased the power of the Darrat and Zhong (2005) analysis, but there is also the suggestion that perhaps the linkages were gradually strengthening over time.

Our empirical study of information linkages between the three NAFTA countries adds to the evidence. We analyse daily spot return data on the three country equity indices over the period January 1990–October 2005, and endeavour to demonstrate the existence of strong information linkages in a

¹Hufbauer and Schott (2005) provide a good overview of developments up to 2005.

direct way, as opposed to inferring them from secondary macroeconomic variables. The novel aspect of our analysis is the use of an appropriate stochastic volatility model which allows us to disentangle the complicated associations in the data by a direct analysis of market returns. The data integrity is subjected to rigorous checks and the robustness of our results is verified in a number of ways, including by simulations of the underlying stochastic model, and verification of distributional assumptions underlying the model.

An understanding of equity market correlations and volatility linkages is important to assess possible benefits of diversification across markets within one economy or across countries, and concomitant benefits of market-wide or country-wide hedging strategies. If volatility linkages are not estimated with precision, hedging strategies may be suboptimal or even counterproductive. Our analysis allows us to quantify these aspects of the relationships between the three NAFTA countries, before and after the inception of the NAFTA.

The paper is organised as follows. Some background, both to the NAFTA itself, and to the models we use, are in Section 2. The data is described in Section 3, and the methods used to analyse them in Section 4. Section 5 gives the main model-fitting results, while Section 6 reviews the appropriateness of the methodology used, with particular reference to robustness issues. Section 7 draws together the findings and considers their implications, and Section 8 concludes the discussion. In an [Appendix](#), we derive an expression for the (rather intractable) likelihood of the model.

2 Background

NAFTA evolved slowly over many years and came finally into effect on January 1, 1994. It is said to be the most comprehensive free trade pact ever negotiated between regional trading partners (Hufbauer and Schott 1993). It sets out specific rules for the imposition of trade restrictions for every good and service traded among the three participating nations, and attempts to promote investment among the three countries.

The objective of any free trade agreement is to lower or eliminate barriers to trade and associated capital controls. Ideally, in economic terms, this increases the welfare of citizens in the signature countries of the agreement. However, it is extremely difficult to quantify or measure this increase in welfare as these effects are not immediate and are not necessarily symmetric for the participating markets. The observations of Gould (1996, 1998), documenting increased trade activity among the NAFTA partners, and of Stiglitz (2004), reporting a failure to reduce income disparities between the US and Mexico, do not necessarily imply the success or failure of these aspects of NAFTA, as one does not know how trade activities or income disparities would have turned out without the agreement. Perhaps, in relative terms, they may in fact have been reduced. Currently, the evidence is conflicting as to the success or otherwise of NAFTA, which highlights how difficult it is to measure the benefits and success of a trade agreement in broad economic terms.

From a financial perspective, however, these issues are a little clearer. It is possible to make stronger predictions about what effect the decrease in barriers of trade and capital restrictions will have on the market participants. The effect of a free trade agreement, if there is one, should be a growth in the common information available to the markets involved, and in the information transmitted between them. Our aim is to examine this issue in relation to NAFTA by means of a stochastic volatility model which allows us to concentrate on associations between volatilities of returns as well as on returns correlations themselves. Both kinds of associations may be expected to increase as a result of a free trade agreement, and the model allows for this.

As we shall see in Section 3, a simplistic descriptive analysis of the data does not detect significant changes in volatility linkages after the introduction of NAFTA. But simple data summaries are subject to many confounding effects² and a more sophisticated analysis is needed to disentangle the complicated associations in the data. Details of this are in Section 5. We now provide some background to the approach we decided to adopt.

The stochastic volatility (SV) model we use is based on a stochastic information-transfer paradigm, which was initiated by Taylor (1986). In this model, the variance of the squared log returns is regarded as an unobserved, latent variable, which depends on random noise in addition to past conditional variance. It was further developed by Taylor (1994), Ruiz (1994), and Harvey et al. (1994). Under a rational expectations paradigm, volatility linkages are assumed generated from two sources: firstly, from common information that affects all markets; and secondly from information spillovers consequent on cross-market hedging. Originally, it was assumed in the SV model that information events are independent of one another, and that trading volumes are approximately normally distributed, assumptions which are not supported empirically. The model was subsequently extended by modelling the daily information events as an autoregressive process; Lamoureux and Lastrapes (1990), Andersen (1996), Fleming et al. (1998). A further significant extension of the model, essentially into the form we use here (see Section 4.1), was to bivariate or multivariate situations.

The model, in its original or an extended form, has been used extensively; for example, Fleming et al. (1998) applied a bivariate version of it to assess

²These are of course numerous. We quote from Kose et al. (2005) “Isolating the effects of NAFTA on its partner countries is particularly difficult given the significant other shocks that have occurred over the past decade. NAFTA undoubtedly had a significant impact on the macroeconomic environment facing Mexico . . . (but) distinguishing the effect of these changes is complicated by the fact that many . . . were anticipated well in advance of the agreement’s ratification, and by the fact that the liberalization was phased in only gradually. In addition, a host of significant other ‘shocks’ had important effects on Mexico and its NAFTA partners during this period, including: (i) the severe financial crisis that Mexico suffered in 1994 (the Tequila crisis), which forced a sharp devaluation of the peso; (ii) the wide range of other free trade arrangements that the NAFTA partners signed during the same period; and (iii) the broader global cyclical environment, which included a recovery from recessions in the early 1990s, the boom through to the end of the decade, and the more recent global slump”.

market volatility linkages in US markets using a generalized method of moments (GMM) framework.

In preliminary analysis using simulation studies, briefly reported in Section 7.4, we found that neither GMM nor quasi maximum likelihood (QML) methods provided reliable estimates of the parameters in the SV model, which led us to explore the use of Monte Carlo Markov Chain (MCMC) methods to fit it. Several studies, in a related context, such as those of Jacquier et al. (1994) and Broto and Ruiz (2004), have suggested that MCMC is a more efficient and informative estimation methodology than the others.

Hence, using MCMC methods, we present a systematic and rigorous investigation of market linkages via associations between the market volatilities which reflect only information transfers. The methodology extends readily to other related situations such as capital markets integration. A stable and robust methodology to investigate such market linkages is important as increased information exchange among countries that participate in any trade agreement is the ultimate goal. This appears to have been achieved through NAFTA. We expand on this further in the Discussion, Subsection 7.1.

3 Data

Our analysis is based mainly on equity market returns (filtered spot data) for the US, Mexico and Canada. These were proxied by appropriate indices extracted from Thomson Datastream International. The approach is essentially event study based. To cover the NAFTA period, pre and post, we took daily data extending from January 1990, 4 years prior to the introduction, to October 2005, 10 years after.³ For the US return index, we used S&P500 data, for Canada the S&P/TSX Composite Index, and for Mexico the Mexican Data Stream Market Index. The Mexican peso and Canadian dollar values were converted into US dollars using daily exchange rates.

3.1 Filtering US & Canadian spot data for non-synchronicity

Previous studies on volatility linkages, such as that of Fleming et al. (1998), have had the advantage of using futures data rather than spot data, thereby avoiding the problem of non-synchronous trading inherent in spot indices. As discussed in Ahn et al. (2002), some nonsynchronicity will be induced into the stock index because the stock index value is recorded at the end of the trading day using the last transaction price of the day, although shares do not trade exactly at the end of the day. The nonsynchronicity problem is minimized by the use of futures data, as futures are more liquid, and summarize in a single

³The US fiscal year ends in October. G.W. Bush was elected US President in November 2004. He supported NAFTA in various ways. We allowed a year following his inauguration for changes to work through.

contract the market movements of the day. For the US and Canadian markets futures data is available for the complete period in this study, whereas for Mexico this is not the case, as the Mexican futures market did not exist for the entirety of the period of interest. Consequently our main analysis is based on spot data, which is available for the Mexican market over the entire period of interest, together with the corresponding spot data for the US and Canadian markets.

To deal with the possible problem of non-synchronous prices, we followed a process developed by Stoll and Whaley (1990), who suggested the use of an ARMA model to filter spot data and remove distortions caused by non-synchronous trading. In order to test whether the filter successfully dealt with the nonsynchronicity problem, we conducted a preliminary analysis, discussed in Section 6.5 below, using a bivariate version of the SV model fitted separately to the US–Canada spot and futures data. In that analysis, we found no significant differences between the two sets of parameter estimates, leading us to conclude that the ARMA filters performed as required, in removing any nonsynchronicity in the spot series. Consequently, we restrict our attention to the spot series for all three countries for the remainder of the paper. We turn now to describing this data, and the analysis to be done on it.

3.2 Descriptive statistics of spot data

Table 1 displays some sample statistics of the spot return data for the three equity series, calculated over three periods. The complete period runs from 1 January 1990 to 21 October 2005 and contains 4124 daily observations. The first sub-period is the pre-NAFTA period which ends on 31 December 1993 and consists of 1044 daily observations. The post-NAFTA period starts on 1

Table 1 Descriptive statistics for the US, Canadian and Mexican equity markets spot data

	Returns (p.a.)	Std.Dev. (p.a.)	Skewness	Excess Kurtosis	Autocorr of squared rets
USA					
Full period	0.0943	0.17	−0.11 (0.01)	3.98 (0.00)	0.20 (13.0)
Pre-NAFTA	0.0975	0.12	−0.05 (0.51)	2.27 (0.00)	0.11 (4.0)
Post-NAFTA	0.0933	0.17	−0.11 (0.01)	3.76 (0.00)	0.20 (11.3)
Canada					
Full Period	0.0779	0.15	−0.68 (0.00)	6.49 (0.00)	0.12 (8.0)
Pre-NAFTA	0.0210	0.09	−0.21 (0.01)	1.11 (0.00)	0.07 (2.4)
Post-NAFTA	0.0972	0.17	−0.68 (0.00)	5.50 (0.00)	0.11 (5.9)
Mexico					
Full period	0.1385	0.27	−0.43 (0.00)	16.20 (0.00)	0.29 (20.2)
Pre-NAFTA	0.4244	0.21	0.09 (0.24)	4.32 (0.00)	0.11 (3.9)
Post-NAFTA	0.0416	0.29	−0.47 (0.00)	16.37 (0.00)	0.31 (17.9)

Reported are the sample mean and standard deviation of annualised daily log returns, skewness, excess kurtosis (p -values against a null of zero skewness and normal kurtosis are in parentheses), and the first order autocorrelation of squared returns (t -statistic in parentheses). The daily returns were annualised on the basis of 252 trading days per year

Table 2 Correlations of returns across the equity markets of the US, Canada and Mexico (*t*-statistic in parentheses)

	Complete period			Pre-NAFTA			Post-NAFTA		
	US	Canada	Mexico	US	Canada	Mexico	US	Canada	Mexico
US	1	0.63 (52.6)	0.40 (28.6)	1	0.53 (21.6)	0.25 (8.8)	1	0.65 (47.0)	0.44 (26.9)
Canada		1	0.38 (26.5)		1	0.14 (5.0)		1	0.42 (25.6)
Mexico			1			1			1

January 1994 and runs through until 21 October 2005, consisting of 3080 daily observations. As the table reveals, changes in the descriptive statistics from the pre- to the post-period are not at all consistent across the three countries. In the US, returns decreased slightly after the introduction of the agreement, whereas in Canada, returns increased from 2.10% to 9.72% per annum. Mexico suffered a big reduction in returns, with a pre-NAFTA return of 42.44% per annum, which decreased to 4.16% per annum, post-NAFTA. Standard deviations increased for all three NAFTA participants. In all cases, the returns are significantly negatively skewed and leptokurtic, as commonly found in high frequency financial data. (All Jacque-Bera statistics are significant at the 0.1% significance level.) Before further analysis, the data was de-meant and filtered, as described in the previous subsection.

Table 2 reports the correlations of returns across the three equity markets for the different periods. It shows that there are existing relationships, which are significant in all cases, in the pre-NAFTA period. These relationships strengthened in the post-NAFTA period, especially between the US and Mexico and between Canada and Mexico.

Inspection of Table 1 shows that the squared returns are significantly autocorrelated in all periods for each of the three countries, suggesting the presence of some persistence of volatility in the data, and pointing to the need for a stochastic volatility model to accommodate it.

Table 3 shows the correlations of squared returns across the equity markets of the US, Canada and Mexico for the different periods. In extant literature, returns squared have often been used as an approximation for volatility, hence, their correlation as a proxy for volatility correlations. From this table, we see that the squared returns are significantly correlated between the three countries in all periods, except for Canada and Mexico, pre-NAFTA. Further,

Table 3 Correlations of squared returns across the equity markets of the US, Canada and Mexico (*t*-statistic in parentheses)

	Complete period			Pre-NAFTA			Post-NAFTA		
	US	Canada	Mexico	US	Canada	Mexico	US	Canada	Mexico
US	1	0.53 (40.8)	0.21 (13.9)	1	0.54 (22.2)	0.12 (4.3)	1	0.54 (35.4)	0.22 (12.2)
Canada		1	0.20 (13.5)		1	0.21 (0.7)		1	0.21 (11.8)
Mexico			1			1			1

there is little evidence of changes in those correlations from the pre- to the post-NAFTA periods, except perhaps for the US-Mexico correlation. But these relatively crude data summaries are subject to many confounding effects and a much more detailed modelling exercise is need to disentangle the complicated associations in the data. We turn to this analysis next.

4 Methods

In Section 4.1 we briefly outline the precise version of the model to be used in the analysis, and Section 4.2 reviews the Bayesian MCMC approach needed to fit the model.

4.1 A trivariate stochastic volatility model

A number of multivariate SV models have been proposed in the literature, but the one to be used in this paper has the advantage of being directly based on an information transfer paradigm. It is an N -variate extension of the stationary SV model as given in Taylor (1994) and Ruiz (1994), where in our case the observations, y_{it} , are taken to be the mean corrected (ARMA-filtered) log returns of the i -th country's stock market index, $i = 1, 2, \dots, N$, at time t , while the log squared volatilities follow AR(1) processes. In the NAFTA case, $N = 3$ since we have three participating countries, so we have a trivariate stochastic volatility model, which, together with several sub-models, can be subsumed within the following two model equations. First, assume

$$\mathbf{y}_t = \mathbf{H}_t^{1/2} \boldsymbol{\varepsilon}_t, \quad (1)$$

where $\mathbf{y}_t = (y_{1t}, y_{2t}, y_{3t})$, and the subscript t indicates daily observations over a time period $t = 1, 2, \dots, T$. The matrix $\mathbf{H}_t = \text{diag}(e^{h_{1t}}, e^{h_{2t}}, e^{h_{3t}})$ is a 3×3 diagonal matrix whose diagonal log-elements are assumed to satisfy the recursion

$$\mathbf{h}_t = \boldsymbol{\mu} + \boldsymbol{\Phi}(\mathbf{h}_{t-1} - \boldsymbol{\mu}) + \boldsymbol{\eta}_t, \quad t = 1, 2, \dots, T, \quad (2)$$

where $\mathbf{h}_t = (h_{1t}, h_{2t}, h_{3t})$, and \mathbf{h}_0 is a given vector. The matrix $\boldsymbol{\Phi} = (\varphi_{ij})_{i,j \in \{1,2,3\}}$ parameterises a lagged regression of log-volatilities on previous log-volatilities, allowing for a possible cross-country lagged relationship. Long-term mean levels of the log-volatilities are captured by the vector $\boldsymbol{\mu}$.

The distributions of the random quantities $\boldsymbol{\varepsilon}_t$ and $\boldsymbol{\eta}_t$ are assumed to satisfy:

1. the N -vectors $\boldsymbol{\varepsilon}_t = (\varepsilon_{1t}, \dots, \varepsilon_{Nt})$ are independent and identically distributed (i.i.d.) over $t = 1, 2, \dots, T$, having the same distribution as a generic vector $\boldsymbol{\varepsilon} = (\varepsilon_1, \dots, \varepsilon_N)$, which furthermore has mean vector 0 and finite covariance matrix $\boldsymbol{\Sigma}_\varepsilon$, of which the diagonal elements are all equal to 1;
2. the N -vectors $\boldsymbol{\eta}_t = (\eta_{1t}, \dots, \eta_{Nt})$ are i.i.d. over $t = 1, 2, \dots, T$, having the same distribution as a generic vector $\boldsymbol{\eta} = (\eta_1, \dots, \eta_N)$, which furthermore has mean vector 0 and finite covariance matrix $\boldsymbol{\Sigma}_\eta$;

- the collections $\{\varepsilon_{it}; i = 1, 2, \dots, N, t = 1, 2, \dots, T\}$ and $\{\eta_{it}; i = 1, 2, \dots, N, t = 1, 2, \dots, T\}$ are independent of each other.

The correlation matrix Σ_ε quantifies the correlations of the log-returns, standardised for the dynamic stochastic volatility, across the three countries, while the matrix Σ_η measures the variances and cross-country covariances in the error terms of the log-volatilities. These matrices are assumed to have the forms:

$$\Sigma_\varepsilon = \begin{pmatrix} 1 & \rho_{\varepsilon 12} & \rho_{\varepsilon 13} \\ \rho_{\varepsilon 21} & 1 & \rho_{\varepsilon 23} \\ \rho_{\varepsilon 31} & \rho_{\varepsilon 32} & 1 \end{pmatrix} \quad \text{and} \quad \Sigma_\eta = \begin{pmatrix} \sigma_1^2 & \sigma_{12} & \sigma_{13} \\ \sigma_{21} & \sigma_2^2 & \sigma_{23} \\ \sigma_{31} & \sigma_{32} & \sigma_3^2 \end{pmatrix}.$$

We will write $\sigma_{ij} = \rho_{\eta ij}\sigma_i\sigma_j$ for $i, j = 1, 2, 3$ and report the estimates of $\rho_{\eta ij}$ rather than the estimates of σ_{ij} , thus, the correlations rather than the covariances of the log volatilities. Note that Σ_ε and Σ_η are symmetric matrices whereas the matrix Φ is not necessarily symmetric, reflecting possible asymmetric transmission of information between countries.

So far we have not made any distributional assumptions concerning ε and η , other than the finiteness of their second moments. Most of our analysis will assume normal distributions for ε and η , but we will also consider multivariate t -distributions for them, in a robustness study reported in Section 6.4.

4.2 Bayesian inference and the MCMC method

It is possible to write down the likelihood corresponding to the model specified by Eqs. 1 and 2 in a more-or-less explicit form, assuming normality of the residual terms, as we do in the Appendix, but the resulting expression is a mixture distribution (as seen in Tauchen and Pitts (1983) in the univariate case) requiring integration over a high dimensional space (in our case, of dimension T , the number of time observations), and is non-tractable by usual means. (See Eq. 7 of the Appendix.) But its mixture/latent variable aspect is very well suited to the use of MCMC methods (Chib et al. 2002), which were developed to deal with this kind of situation, and we chose to investigate these. To be more specific, we used the Gibbs sampling algorithm (see Gilks et al. (1996) and Robert and Casella (2000)) to fit the SV model specified by Eqs. 1 and 2, together with various interesting sub-models.

The fundamental idea behind the MCMC method in Bayesian applications, such as this one, is to produce variates from their posterior distribution by repeatedly sampling a Markov chain whose invariant distribution is the target distribution of interest. We used the WinBUGS 1.4.1 program, which is a freely available Bayesian software (Spiegelhalter et al. 2003), to fit these models in the Bayesian context. WinBUGS is an object-oriented program, so the user does not have access to all details and is only informed about the interfaces. To reassure ourselves as to its reliability for our application, we conducted

a simulation study (cf. Section 6.3) in which WinBUGS proved to be very accurate for our type of model.⁴

To apply the MCMC method in our situation, we need to specify prior distributions for each of the following parameters in the model: $\rho_{\varepsilon ij}$, μ_i , φ_{ij} , σ_i , and $\rho_{\eta ij}$ ($i, j = 1, 2, 3$). Prior independence of all these parameters is assumed and all prior marginal distributions were proper. For the correlations $\rho_{\varepsilon 12}$, $\rho_{\varepsilon 13}$, $\rho_{\varepsilon 23}$, $\rho_{\eta 12}$, $\rho_{\eta 13}$ and $\rho_{\eta 23}$, we used a uniform prior on $[-1, 1]$ (therefore uninformative), while the components of the mean vector μ were assumed to have normal priors with mean -9 and standard deviation 5, each. This was based on the fact that the descriptive statistics in Table 1 indicate an average log-volatility of approximately -9 (see Section 3.2). A standard deviation of 5 is chosen so that these priors are quite uninformative. Further, while we might expect the standard deviations σ_i to be rather small, we did not want to impose any further restrictions on their prior distributions. Hence, we simply chose uniform priors on $[0, 1]$ for the σ_i , which subsequently proved to have a sufficiently large support. We took a Beta(20, 1.5) prior for the transformed parameter $(\varphi_{ii} + 1)/2$ for each $i = 1, 2, 3$, so that φ_{ii} ($i = 1, 2, 3$) has support $[-1, 1]$, a prior mean of 0.8605 and a prior standard deviation of 0.1074. This reflects the high degree of persistence in volatility commonly found for weekly or daily data. In a robustness study, reported on in Section 6.2, we varied those prior distributions to assess the effect on the parameter estimates. Finally, for the parameters φ_{ij} ($i \neq j$) we took a uniform prior on $[-1, 1]$.⁵

The model described in the previous section is the main model of interest in this paper; however, in line with standard statistical procedures, we also considered several interesting sub-models. These were derived by equating some parameters, such as the cross-country correlations, to zero.

The aim of the analysis is to select a model containing those parameters, and only those parameters, which are necessary to adequately describe the data. To get a feel for which parameters are significant, we proceeded initially by fitting various sub-models to the (larger) post-NAFTA data set. This procedure also allowed us to build up information on good starting values for the WinBUGS program. We started with a submodel where all elements in the Φ matrix were estimated, but only the diagonal elements of the Σ_{η} matrix. In this setup, the 95% credible intervals for parameters φ_{13} , φ_{31} , φ_{23} , φ_{32} , contained

⁴For details on the MCMC procedure used by WinBUGS, we refer to Lunn et al. (2000). In particular, this paper describes the architecture of the six primary subsystems of WinBUGS: Graph, Updater, Monitors, Bugs, Samples, Doodle. The updater objects are responsible for carrying out the MCMC simulation. There are updaters for specific distributions, i.e., when a full conditional distribution can be expressed in closed form, and general updaters, such as for Metropolis-Hastings samplers.

⁵We ran the WinBUGS sampler for a total of 15000 iterations and discarded the first 5000 iterations for burn-in. These numbers were decided on after conducting simulation trials with parameters close to those found for the spot data. The chains always converged within 3000 to 4000 iterations, and gave very consistent estimates from 10000 iterations after this initial burn-in period. We found it to be very important to choose appropriate starting values for the series of log-volatilities, which we obtained by running the program for simpler submodels.

zero, so in the next step they were omitted from the model, and instead, correlation parameters were included in the matrix Σ_η . Again, considering the 95% credible intervals for this sub-model showed that the off-diagonal elements of Σ_η were necessary, but that the parameters φ_{12} and φ_{21} were not. The final model, described in detail in Section 4.1, thus contains estimates of all elements in Σ_η and Σ_ε , and of course of μ , but only for the diagonal elements of Φ .

5 Main results

Table 4 shows the posterior mean estimates obtained by fitting the trivariate SV model to the de-meaned, filtered returns, for all three countries. The 95% credible intervals are also reported in Table 4. All parameters are significantly different from zero, for both the pre- and post-NAFTA data, as judged by the credible intervals. (As expected, these intervals are wider for the pre- than for the post-period, since the pre-period data set contains only around one third as many as observations as the post-period data set.)

The estimates of μ_1 , μ_2 and μ_3 in Table 4 represent the long-term mean levels of the log-volatilities.⁶ This level has increased, significantly in the case of Canada and Mexico, from the pre- to the post-NAFTA period in all cases. The estimates of σ_1 , σ_2 and σ_3 are of the long-term standard deviations of the log-volatilities. These standard deviations are reduced (significantly in the case of Canada and Mexico) in all three countries in the post-NAFTA period. This suggests a stabilizing effect of NAFTA. The estimates of φ_{11} , φ_{22} and φ_{33} represent a lagged regression of log-volatilities on previous log-volatilities. These have all increased from the pre- to the post-NAFTA period and are much more similar now from country to country, again suggesting a stabilizing effect of NAFTA. Volatility appears to be stationary, with autocorrelations at most around 0.98. The estimates of $\rho_{\varepsilon_{12}}$, $\rho_{\varepsilon_{13}}$ and $\rho_{\varepsilon_{23}}$ measure the cross-country correlations of returns in our model. They are quite comparable with those in Table 2.

The estimates of $\rho_{\eta_{12}}$, $\rho_{\eta_{13}}$ and $\rho_{\eta_{23}}$ represent the main parameters of focus for this paper. They quantify the volatility linkages, representing information linkages according to the Fleming et al. (1998) paradigm. These are significantly positive in all cases and have increased greatly from the pre- to the post-NAFTA period, especially in the case of the US and Canada and Canada and Mexico. This indicates that information flows are greater and equity markets are more closely linked for all three countries after the inception of the free trade agreement. This conclusion is quite different from that which would be

⁶As a check, note that Eq. 1 implies $Ey_{it}^2 = e^{\mu_i + \sigma_i^2/2}$, for each $i = 1, 2, 3$, and $t = 1, 2, \dots, T$. Thus overall volatility as per the model is $\sqrt{252}e^{\mu_i + \sigma_i^2/2}$, on an annualised basis. Substituting the estimates $\hat{\mu}_1 = -9.744$, $\hat{\sigma}_1 = 0.1447$, etc., we find a close correspondence with the volatilities reported in Table 1 (0.12, 0.08, etc.).

Table 4 Posterior mean estimates and 95% credible intervals for the parameters in the full model

Parameters	Pre-period		Post-period	
	Post. mean	95% Cred. int.	Post. mean	95% Cred. int.
μ_1	-9.7440	[-10.120, -9.3340]	-9.4672	[-9.7910, -9.2210]
μ_2	-10.4600	[-10.630, -10.280]	-9.4319	[-9.7240, -9.1480]
μ_3	-8.9350	[-9.1190, -8.7340]	-8.5260	[-8.8460, -8.1879]
σ_1	0.1447	[0.0978, 0.1896]	0.1339	[0.1188, 0.1509]
σ_2	0.3304	[0.2341, 0.4296]	0.1835	[0.1667, 0.1996]
σ_3	0.4959	[0.3881, 0.6166]	0.2659	[0.2306, 0.3036]
φ_{11}	0.9795	[0.9650, 0.9934]	0.9857	[0.9798, 0.9909]
φ_{22}	0.8676	[0.7893, 0.9404]	0.9783	[0.9699, 0.9859]
φ_{33}	0.7855	[0.6933, 0.8790]	0.9697	[0.9568, 0.9827]
$\rho_{\varepsilon 12}$	0.5063	[0.4573, 0.5514]	0.6202	[0.5979, 0.6417]
$\rho_{\varepsilon 13}$	0.2872	[0.2272, 0.3459]	0.5041	[0.4780, 0.5295]
$\rho_{\varepsilon 23}$	0.1992	[0.1349, 0.2627]	0.4587	[0.4304, 0.4869]
$\rho_{\eta 12}$	0.2669	[0.0223, 0.5039]	0.7895	[0.7316, 0.8404]
$\rho_{\eta 13}$	0.3685	[0.0603, 0.6482]	0.4603	[0.2910, 0.6252]
$\rho_{\eta 23}$	0.2512	[0.0104, 0.5008]	0.5731	[0.4210, 0.7438]

Subscript 1 denotes the US, subscript 2 denotes Canada and subscript 3 denotes Mexico. 1044 daily returns in the pre-period and 3080 in the post period

drawn from inspection of Table 3. Here, we see that the SV methodology highlights linkages of returns which are not apparent with a simplistic analysis.

Further, the SV methodology reveals strong linkages existing *pre*-NAFTA, which we should expect, as there was a prior existing free-trade agreement between the US and Canada (the Canada–US Free Trade agreement, CUFTA). Mexican firms have long used the US debt market for financing their operations (Lederman et al. 2003), hence, we would also expect pre-existing information linkages for these countries as well. The analysis has picked up these pre-existing relationships as well as a significant change from the pre-to post NAFTA period, which was obscured by the simple proxy comparison in Table 3.

To illustrate the precision of the parameter estimates in Table 4 and highlight the most significant changes caused by NAFTA, Fig. 1 contains density plots of the estimated marginal posterior distributions for the most important model parameters. Each of the nine plots shows the estimated densities for both the pre- and the post-NAFTA period. In the first row are the volatilities σ_1 , σ_2 , and σ_3 in the log-volatility Eq. 2 of the SV model. The second row shows densities for the return correlations $\rho_{\varepsilon 12}$, $\rho_{\varepsilon 13}$, and $\rho_{\varepsilon 23}$ in Eq. 1 of the trivariate SV model. The strong increase in all three correlations after NAFTA came into effect is clearly seen. Interestingly, here the US market experienced the smallest change, as might be expected since the US market is the most dominant one among these three countries. However, the volatility of the log-volatility in the Canadian and the Mexican market decreased quite significantly. In the third row, the log-volatility correlation $\rho_{\eta 12}$ between the US and Canada increased the most, whereas the volatility correlations $\rho_{\eta 13}$ and $\rho_{\eta 23}$ between the US and Mexico and between Canada and Mexico, respectively, show only a slight tendency to be stronger since 1994.

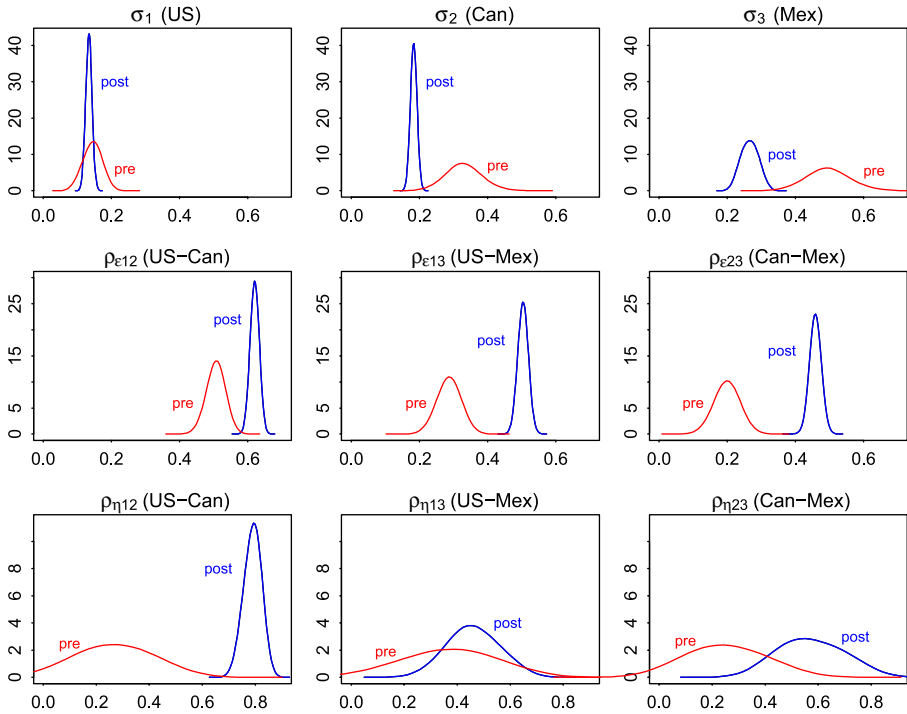


Fig. 1 Estimated marginal posterior density plots, pre- and post-NAFTA. *First row:* volatilities σ_1 , σ_2 , σ_3 . *Second row:* return correlations $\rho_{\epsilon 12}$, $\rho_{\epsilon 13}$, $\rho_{\epsilon 23}$. *Third row:* log-volatility correlations $\rho_{\eta 12}$, $\rho_{\eta 13}$, $\rho_{\eta 23}$

6 Robustness checks

In this section we outline some extra sets of analyses we carried out on the data, in order to assess the goodness-of-fit of the SV model to the NAFTA data (Section 6.1), to assess (in Section 6.2) the sensitivity of the results to the assumptions made concerning prior distributions in Section 4.2, to examine the reliability of the estimates, which we did by simulation in Section 6.3, and to examine the robustness of the model to the normality assumption for the returns in Section 6.4. Section 6.5 contains the spot versus futures comparative analyses which validate the use of the spot data after filtering to reduce the effects of nonsynchronicity.

6.1 Goodness-of-fit

The MCMC method is based on a parametric model description of the data and hence can be expected to perform much better than estimating-equation methods such as, for example, the GMM method. But ultimately the model’s appropriateness depends on its goodness of fit to the data; is it an adequate

descriptor? For assessing goodness of fit of a model to the observed data within a Bayesian framework, a suitable approach is a *posterior predictive analysis* (PPA). This kind of analysis has been thoroughly studied in the literature; see for example Gelman et al. (1995) or Gelman et al. (1996).

The PPA requires the estimation of the *posterior predictive distribution*. Samples from this distribution can be obtained from the MCMC procedure. We did this for the post-NAFTA data, using the iterations numbered $i = 5010, 5020, 5030, \dots, 15000$ from the WinBUGS MCMC output. Thus the analysis is based on 1000 replicated data sets each of length 3080. The discrepancy between the chosen model and the observed data is investigated by calculating the value of a discrepancy function which measures the difference between the replicated data sets and the original data.

As pointed out by Gelman et al. (1996), the choice of discrepancy function should reflect the inferential interests for the problem at hand. After deciding on these, the discrepancy statistics of the replicated data sets are compared with those of the observed data set. Since we are most interested in assessing the correlations between, and the volatilities within each of the three NAFTA markets separately, we computed for each i the following six discrepancy statistics:

$$D_{(j)}(\mathbf{y}, \boldsymbol{\theta}_i) := D_{(j)}(\mathbf{y}) := \sum_{t=1}^T y_{jt}^2, \quad j = 1, 2, 3,$$

$$D_{(j,k)}(\mathbf{y}, \boldsymbol{\theta}_i) := \widehat{\text{covr}}((y_{j1}e^{-h_{j1}/2}, \dots, y_{jT}e^{-h_{jT}/2}), (y_{k1}e^{-h_{k1}/2}, \dots, y_{kT}e^{-h_{kT}/2})),$$

$$(j, k) \in \{(1, 2), (1, 3), (2, 3)\}.$$

Here $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_T)$ represents the original data, and we let \mathbf{y}^{rep} represent the replicated data sets. $\boldsymbol{\theta}_i$ collects all parameters in the trivariate SV model, including the latent volatility process. Note that the statistics $D_{(j)}$, $j = 1, 2, 3$, do not depend directly on $\boldsymbol{\theta}_i$, whereas $D_{(1,2)}$, $D_{(1,3)}$, and $D_{(2,3)}$ do, since they depend on h_{j1}, h_{j2}, \dots . The replicated data sets, of course, always depend on $\boldsymbol{\theta}_i$.

The 1000 values of $D_{(j)}(\mathbf{y}^{\text{rep}})$ from the replicated data sets are summarized for each $j = 1, 2, 3$, in the histograms in the first row of Fig. 2. For comparison, the values of $D_{(j)}(\mathbf{y})$, $j = 1, 2, 3$, representing the discrepancy statistics for the post-NAFTA data, are added as vertical lines. Since these vertical lines lie well in the middle of the support of the histograms, in fact they are quite typical values, we conclude that there is no significant evidence against the model.

Similarly, we get 1000 values for the discrepancies $D_{(j,k)}$, also from the post-NAFTA data set. In the second row of Fig. 2 are the corresponding histograms for $D_{(j,k)}(\mathbf{y}^{\text{rep}}, \cdot)$, with vertical solid lines representing the mean of $D_{(j,k)}(\mathbf{y}, \boldsymbol{\theta}_i)$ and vertical dashed lines representing the corresponding 2.5% and 97.5% quantiles. Therefore, as in the first row, the histograms are reflecting the discrepancy statistics for the replicated data, whereas the vertical (solid and dashed) lines refer to the discrepancy statistics derived for the post-NAFTA

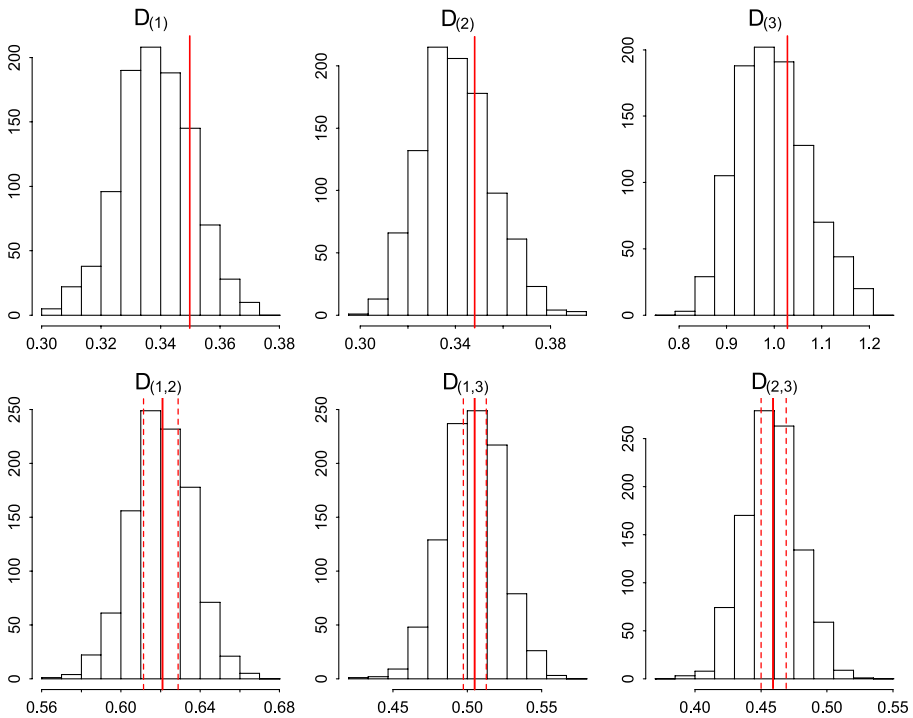


Fig. 2 Posterior predictive analysis. Histograms of discrepancy statistics for replicated data. Vertical (solid and dashed) lines refer to the discrepancy statistics derived for the post-NAFTA data

data. Again the diagrams indicate that our model describes the correlation structure of the data very well.

6.2 Sensitivity to priors analysis

The MCMC method requires the specification of prior distributions for the parameters, as was done in Section 4.2. To assess the sensitivity of the estimation to changes in these, we varied the choice of hyperparameters and distributions for the priors of each parameter. In particular, we wished to check whether the Beta-priors assumed for φ_{ii} , $i = 1, 2, 3$, although being well motivated by the commonly found high persistence in volatility, have too strong an impact on the posterior distribution. To this end, we re-estimated the pre- and post-NAFTA data sets, with the Beta-priors for the φ_{ii} replaced by uniform priors on $[-1, 1]$. The results can be found in Table 5, which can be compared with those from the MCMC run with the Beta-priors in Table 4.

The changes in the estimates with the uniform prior are negligible, and we conclude that priors as disparate as the beta or uniform on the φ_{ii} have little or no effect on the posterior inference. In another analysis, when the variance of the normal priors on the μ_i was increased, similar results were obtained. This

Table 5 Posterior mean estimates and 95% credible intervals for uniform priors on φ_{ii} , $i = 1, 2, 3$

Parameters	Pre-Period		Post-Period	
	Post. mean	95% Cred. int.	Post. mean	95% Cred. int.
μ_1	-9.7467	[-10.123, -9.3365]	-9.4689	[-9.7929, -9.2226]
μ_2	-10.4541	[-10.623, -10.273]	-9.4297	[-9.7223, -9.1451]
μ_3	-8.9361	[-9.1209, -8.7338]	-8.5252	[-8.8455, -8.1869]
σ_1	0.1439	[0.0968, 0.1889]	0.1341	[0.1187, 0.1512]
σ_2	0.3318	[0.2353, 0.4310]	0.1832	[0.1659, 0.2002]
σ_3	0.4981	[0.3901, 0.6195]	0.2670	[0.2311, 0.3048]
φ_{11}	0.9787	[0.9637, 0.9930]	0.9856	[0.9792, 0.9911]
φ_{22}	0.8670	[0.7883, 0.9397]	0.9779	[0.9694, 0.9870]
φ_{33}	0.7857	[0.6929, 0.8800]	0.9694	[0.9569, 0.9822]
$\rho_{\varepsilon 12}$	0.5057	[0.4559, 0.5508]	0.6204	[0.5976, 0.6421]
$\rho_{\varepsilon 13}$	0.2880	[0.2273, 0.3468]	0.5041	[0.4774, 0.5294]
$\rho_{\varepsilon 23}$	0.1988	[0.1339, 0.2621]	0.4590	[0.4307, 0.4874]
$\rho_{\eta 12}$	0.2684	[0.0262, 0.5034]	0.7887	[0.7300, 0.8398]
$\rho_{\eta 13}$	0.3697	[0.0573, 0.6510]	0.4591	[0.2891, 0.6263]
$\rho_{\eta 23}$	0.2498	[0.0116, 0.4985]	0.5725	[0.4193, 0.7420]

Subscript 1 denotes the US, subscript 2 denotes Canada and subscript 3 denotes Mexico

robustness to changes in the prior distributions could be expected since we have a substantial amount of data available, both in the pre- and post-periods, but the extra analyses are reassuring in this respect.

6.3 Simulations

In order to test the accuracy of the MCMC method for our specific model, we ran two sets of simulations. In the first set, we took 1000 observations, approximately the size of pre-NAFTA period data, generated according to the model described in Eqs. 1 and 2, with “true” parameter values close to the fitted values obtained in Table 4, and assuming normal distributions for the ε_t and η_t in Eqs. 1 and 1. The second set consisted of 3000 observations, obtained analogously, to approximate the post-NAFTA period data. We ran the Winbugs analyses on the simulated data to obtain parameter estimates, which can then be compared with the “true” values. The results are reported in Tables 6 and 7.

The number of simulations is rather restricted due to the amount of time needed for each run (approximately 18 h/set for each post-period run), nevertheless it is clear that the model fitting procedure performs very well. It reproduces the true values accurately, with little evidence of bias.

6.4 Fitting the t-distribution

The simulation exercise outlined in Section 6.3, in which we successfully recovered the parameters of the model from simulated data, gives us confidence that MCMC methodology works well on “ideal” data. But in addition to this,

Table 6 Simulation results for datasets of size 1000 using pre-NAFTA period parameter estimates

Parameter	True value	Post. mean	Std. dev.
μ_1	-9.74	-9.70	0.13
μ_2	-10.46	-10.40	0.08
μ_3	-8.93	-8.93	0.08
σ_1	0.15	0.16	0.03
σ_2	0.33	0.32	0.06
σ_3	0.49	0.50	0.06
φ_{11}	0.97	0.95	0.02
φ_{22}	0.87	0.86	0.04
φ_{33}	0.79	0.78	0.04
$\rho_{\varepsilon 12}$	0.50	0.50	0.03
$\rho_{\varepsilon 13}$	0.28	0.29	0.03
$\rho_{\varepsilon 23}$	0.20	0.21	0.04
$\rho_{\eta 12}$	0.26	0.24	0.17
$\rho_{\eta 13}$	0.36	0.31	0.17
$\rho_{\eta 23}$	0.25	0.23	0.11

50 simulations in all were carried out. For each parameter, the table reports the “true” value, the mean of the posterior mean estimates, and the standard deviation of the posterior mean estimates, across the 50 simulations

we carried out a couple of other robustness checks by varying the distributional assumptions.

There is evidence of long-tailedness in the data, as indicated by the significant kurtosis values in Table 1 (skewness, though significant, seems less of a problem as it is small in absolute terms). This suggested fitting a Student *t* rather than a Normal distribution. We did this, using MCMC and the WinBUGS program again, and specifying a model in which the optimal number of degrees of freedom was estimated along with the other parameters of the *t*-distribution. This optimal number turned out to be 12.48 for the pre-NAFTA data set and 13.45 for the post-NAFTA data set. These were rather higher than we expected to find. A *t*-distribution with such high degrees of freedom is very difficult to distinguish, practically, from a normal distribution. Further, the parameter estimates changed little—typically, by less than 10%—from those in Table 4. In particular, the post-NAFTA estimates of $\rho_{\eta 12}$, $\rho_{\eta 13}$ and $\rho_{\eta 23}$ remained high, and highly significantly different from 0. The pre-

Table 7 Simulation results for datasets of size 3000 using post-NAFTA period parameter estimates

Parameter	True value	Post. mean	Std. dev.
μ_1	-9.47	-9.39	0.11
μ_2	-9.43	-9.30	0.10
μ_3	-8.53	-8.48	0.12
σ_1	0.13	0.11	0.01
σ_2	0.18	0.19	0.02
σ_3	0.26	0.26	0.02
φ_{11}	0.98	0.98	0.01
φ_{22}	0.97	0.97	0.01
φ_{33}	0.96	0.96	0.01
$\rho_{\varepsilon 12}$	0.62	0.62	0.01
$\rho_{\varepsilon 13}$	0.50	0.50	0.01
$\rho_{\varepsilon 23}$	0.45	0.45	0.02
$\rho_{\eta 12}$	0.78	0.76	0.05
$\rho_{\eta 13}$	0.46	0.44	0.09
$\rho_{\eta 23}$	0.57	0.56	0.08

50 simulations in all were carried out. For each parameter, the table reports the “true” value, the mean of the posterior mean estimates, and the standard deviation of the posterior mean estimates, across the 50 simulations

NAFTA estimates of these parameters were similar to those in Table 4, but now the credible intervals contained 0 at their lower ends. The autoregressive parameters φ_{ii} moved closer to 1.

As an even more extreme check, we repeated the t -distribution analysis but with degrees of freedom specified to be 4. Such a heavy-tailed distribution is substantially more extreme than is likely to be seen in data of the kind we have. Again, the estimates moved in predictable ways and our general conclusions would not be changed.

On the whole this exercise led us to conclude that the normal distribution parametrization is best for this data set; it appears that allowing for stochastic volatility as we have done compensates to a large extent for the kurtosis in the data. But even taking a t -distribution with degrees of freedom estimated from the data would not change our general conclusions.

6.5 Bivariate model for spot v. futures comparison

To make the filtered spot versus futures comparison for the US and Canada mentioned in Section 3.1, we collected the following futures data. For the Canadian market, we took daily TSE35 contract data (each quarterly contract) from 1 January 1991 until 31 December 1999, after which we took the S&P60 contract data (the S&P60 contract replaced the TSE35 contract at this time) till 21 October 2005. For the US market, we collected daily S&P500 contract data over the entire period, 1 January 1991 to 21 October 2005. In order to generate continuous return series for the futures contracts, the log price relatives for each contract had to be spliced. The splicing was accomplished by selecting the contract with the greatest trading volume so as to decide the date on which to switch from one contract to the next. The day when, on average (from 1 January 1991 until 21 October 2005), the new contract's volume exceeded the previous contract's trading volume, was chosen (Fleming et al. 1998). In the US, the equity futures contract switch-over day is five days, and in Canada it is two days, before the last day of trading of the expiring contract. On the day of the switch, we computed the futures return using the current and previous day's prices for the new contract. On inspection, as expected, the futures returns were found to be smaller than the comparable spot returns, as futures contracts do not include the risk free rate as compensation for the cost of an investment. The dataset includes some extreme periods, such as the technology crash during the year 2000, yielding a large negative annual return. Over the period of analysis, the risk-free rate had substantial variation.

As discussed in Section 3.1, we applied the Stoll and Whaley (1990) technique to the Canadian and US spot data, with the aim of eliminating the nonsynchronicity effects. The bivariate SV model was then fitted using the MCMC method to the filtered spot data and the futures data separately. The lowest Bayesian Information Criterion (BIC) was used to select the optimal ARMA(p,q) specification after fitting all possible combinations of p and q up to a maximum ARMA(4,4) specification. For the US spot equity

Table 8 Posterior mean estimates and 95% credible intervals for the parameters of the spot versus futures comparison, in which the bivariate SV model is fitted separately to filtered spot data (3080 daily returns from the post-NAFTA period) and to futures data

Parameter	Futures data		Filtered spot data	
	Post. mean	95% cred.int.	Post. mean	95% cred.int.
μ_1	-9.5080	[-9.9610, -9.1830]	-9.7460	[-10.3800, -9.3410]
μ_2	-9.3150	[-8.7880, -7.9860]	-9.6080	[-9.9600, -9.3360]
σ_1	0.1625	[0.1196, 0.2041]	0.1145	[0.0903, 0.1361]
σ_2	0.1666	[0.1375, 0.1913]	0.1746	[0.1330, 0.2181]
φ_{11}	0.9815	[0.9719, 0.9911]	0.9891	[0.9830, 0.9950]
φ_{22}	0.9764	[0.9661, 0.9852]	0.9774	[0.9620, 0.9886]
$\rho_{\varepsilon 12}$	0.5791	[0.5567, 0.6015]	0.6185	[0.5966, 0.6401]
$\rho_{\eta 12}$	0.6966	[0.5421, 0.8191]	0.7334	[0.6118, 0.8221]

Subscript 1 denotes the US and subscript 2 denotes Canada. Interpretations of parameter estimates are provided following Table 4

series, this was found to be an ARMA(2,1), and for the Canadian equity series, an ARMA(1,1) model.

Table 8 shows the parameter estimates obtained by fitting the bivariate version of the SV model to the US and Canadian spot data, after passing it through an ARMA filter as described in Section 3.1, compared with the corresponding estimates obtained from futures data. This allows an assessment of whether filtering the spot data had the desired effect of eliminating any synchronicity problem. The bivariate model is simply obtained by selecting $N = 2$ in Eqs. 1 and 2 and making the obvious modifications. The parameter estimates shown in Table 8 are not significantly different between futures and filtered spot, as indicated by the overlapping 95% credibility intervals, hence, we can be reasonably confident that the ARMA filter substantially removes any nonsynchronicity from the spot data.

7 Discussion

7.1 Promotion of information linkages by NAFTA

The SV model initiated by Taylor (1986), as described in Section 2, was preceded by a rational expectations model due to Tauchen and Pitts (1983), in which the market begins each trading day in equilibrium. The economy is assumed to consist of active speculative traders who trade single futures contracts with one another. The traders have different expectations regarding the future; hence they require different transactions in order to satisfy their risk exposures. In the absence of transaction costs, a new equilibrium price will, in principle, be reached after each round of trading following a new information event. Fleming et al. (1998) extended this idea to traders dealing across different markets, postulating that volatility is generated by new

information releases. They argue that investors can be expected to react first to the arrival of new *common information*, that is, to information affecting the different markets simultaneously. Rebalancing portfolios in response to such an information event then causes *information spillovers* via changes in investors' requirements for hedging, for example. Hence, the model predicts that volatility changes across markets which lend themselves to cross-hedging, or, more generally, to cross-trading, are likely to be highly correlated.

The main impediments to perfect correlations of volatility changes across markets are of a practical nature (Fleming et al. 1998), such as the imposition of transaction costs. Also relevant may be capital constraints or other policy constraints. We might expect that in a cross-country setting these practical considerations would be greatly exacerbated. Nevertheless, despite the diluting effects of factors like these, we would still expect very strong volatility linkages across markets, and only those markets, that are driven by the same fundamental macroeconomic factors. The countries we chose to study were selected because the inception of the trade agreement produced an ideal scenario in which to assess the effect of a major, potentially positive, intervention on the flow of financial information between three closely aligned trading partners. Increased information flow between them would mean that information of a macroeconomic nature, such as a change in interest rates, is much more likely to quickly affect a trading partner. Hence, we expected a much higher degree of common information to be available in the post-NAFTA period. Further, following NAFTA, investors are now more likely to cross-hedge across the three markets, thus increasing occurrences of information spillovers from one market to another. Overall, we expect both sources of volatility linkages, common information and information spillovers, to increase from the pre-NAFTA to the post-NAFTA period.

7.2 Significant changes resulting from NAFTA

Using a sound and reliable methodology, we have been able to show that significant volatility linkages do indeed exist between the US, Canadian and Mexican markets, and that these linkages were significantly increased through the introduction of NAFTA. This appears clear from the estimates of $\rho_{\eta 12}$, $\rho_{\eta 13}$ and $\rho_{\eta 23}$ in Table 4 and the density estimates in Fig. 1. The parameter measuring volatility correlation between the US and Canada increased significantly from 0.50 to 0.62, with disjoint 95% credible intervals, from the pre- to the post-period. The change is quite dramatic between the US and Mexican markets; here the volatility correlation increased from 0.29 to 0.51, and, again, considering the credible intervals, this increase is very significant. The same holds for the correlation between Canada and Mexico, which increased from 0.20 before to 0.46 after NAFTA. These volatility correlations are very different from those in Table 3, which lists the values of a common proxy for volatility correlations, namely the correlations of squared returns. On the basis of Table 3, we would naively conclude that there is little evidence of a change in country to country volatility correlations from pre- to post-NAFTA, except

perhaps for a small increase in the case of Mexico and the US. With the SV model and the MCMC methodology, however, a significant increase in all three volatility correlations is revealed. Furthermore, this insight is achieved without sacrificing information on other features of interest in the data, such as the existence of correlations between returns. This highlights the value of fitting a model based on sound financial considerations, over the direct interpretation of marginal effects which may be subject to confounding factors. That a robust methodology for estimating volatility linkages is an important factor when estimating the transmission of information across international equity markets is also consistent with Wongswan (2006), who studied the transmission of information from developed to emerging markets. He finds that fundamental linkages may not be detected due to the nature of information proxies used in extant literature and due to the use of low-frequency data.

7.3 Enhancement of cross-country diversification

An understanding of market linkages is important for cross-country hedging and diversification purposes. Portfolio managers, in particular, are highly sensitive both to levels of volatility and to correlations between asset returns in making investment decisions. But the presence of volatility linkages adds a further level of complexity to asset allocations decisions. The latter cannot result in optimal portfolios if the dependencies between volatilities are ignored, even though correlations between returns may be taken into account. A move to shift funds from a more volatile market to another, apparently less volatile, for example, may be partially nullified by a correlation between the volatilities. This issue is especially acute in the foreign exchange markets, where we may find derivative securities whose payoffs depend on several underlying asset prices, across countries. It is also important to take information on volatility linkages into account when setting appropriate levels of capital adequacy and margin requirements; Chance (1990) reports on the ongoing policy debate on margin requirements for volatility and trading activity.

In a related context, using optimal portfolio considerations, Kodres and Pritsker (2002) give an extension of a rational expectations based model of asset prices due to Grossman and Stiglitz (1980), obtained by formally partitioning the asset values into a component representing the expected value conditional on the information available to informed investors, plus a residual. They use this to further develop the idea by Fleming et al. (1998) of contagion being transmitted between markets by cross market rebalancing, and illustrate with a stylised three country example in which any two of the three countries share a common information factor. Their model has no allowance for volatility persistence though this could be built in.

In general, the effect of disregarding volatility linkages on optimizing portfolios is very difficult to quantify, as, to our knowledge, there is no extension of the Markowitz technique, for example, which takes into consideration such second order effects, though this would seem to be a worthwhile object of study

given our findings.⁷ Any such technology would require as input the estimates that our methodology provides. Nevertheless, by demonstrating the existence and strength of such linkages, as we have done with the NAFTA, the portfolio manager is at least alerted to the need for taking them into account.

7.4 Other methodologies

Previous analyses of the SV model for lower-dimensional data have used GMM and QML methodologies, and, prior to the analyses reported here, we embarked on an extensive analysis of the Tauchen and Pitts (1983) model using the GMM technique employed by Fleming et al. (1998). But we encountered some substantial practical difficulties in specifying and fitting the model. A large number of rather complex moment restrictions greatly complicates the estimation, and we frequently had problems with convergence of the iterations even after the optimization routine was programmed step-by-step in two quite different computer languages (Fortran and the Regression Analysis of Time Series (RATS) statistical package (ESTIMA 2004)). Upon apparent “convergence”, the routines occasionally produced nonsensical results such as correlations outside $[-1, 1]$, etc.⁸ We further found on carrying out simulation exercises that the GMM estimates, when obtainable in simpler models, usually behaved very poorly with non-convergence in approximately 80% of the simulation runs. We found similar with some less extensive analyses using a QML approach. Ultimately we concluded that these methodologies are not as robust and sound as the MCMC methodology, at least with the data and the kind of model we wish to use.

8 Conclusions

It is important to be able to measure in a direct way whether NAFTA or any other trade agreement has aided in the member markets being linked more closely after the inception of the agreement. This paper makes a clear contribution to the literature in this area. It outlines and illustrates the use of a trivariate stochastic volatility model aimed at investigating the existence of, and possible changes in, information linkages between the US, Canada and Mexico, from the pre-NAFTA to the post-NAFTA period. The stochastic volatility model, fitted using MCMC techniques, proved to be very robust when subjected to a number of checks, and to simulations, and generally to provide a more reliable methodology than pre-existing approaches.

With the help of this model we uncovered substantive findings concerning NAFTA, concluding that the lowering or in part elimination of trade barriers is

⁷There are of course many other extensions of Markowitz, allowing for time-varying correlations, skewness in returns distributions, downside risk, etc., but none of the sort mentioned seem to exist.

⁸Lamoureux and Lastrapes (1990) found negative variance estimates in some of their GMM runs.

associated with more common sharing of information across the three markets. When information linkages increase from a pre- to a post-agreement period, we can infer that markets rely upon or trade upon common information, and that information spillovers between the markets are substantial. Our methods and our findings have significant implications also for the assessment of future trade agreements.

Acknowledgements We are grateful for helpful comments and suggestions received from Claudia Czado, Robert Durand, Doug Foster, David Heath, Claudia Klüppelberg, Dilip Madan, and Barry Oliver. Special thanks are due to a referee for penetrating and constructive comments which have resulted in an improved paper.

Appendix: The likelihood

The likelihood for the model in Section 4.1 can be developed as follows, assuming normal distributions for the ϵ_t and η_t . Linearising by taking logs, we set

$$Z_{it} = \log y_{it}^2, \quad \text{and} \quad \xi_{it} = \log \epsilon_{it}^2, \quad \text{for } i = 1, 2, \dots, N, \text{ and } t = 1, 2, \dots, T,$$

where the y_{it} and ϵ_{it} are defined in Eq. 1 and the paragraph following it; recall that $P(\epsilon_{it} = 0) = 0$ for all i and t (and that $N = 3$ in our application). The vectors ϵ_t and η_t are N -vectors with elements $(\epsilon_{it})_{i=1,2,\dots,N}$ and $(\eta_{it})_{i=1,2,\dots,N}$, for $t = 1, 2, \dots, T$, and similarly, let Z_t be the N -vector with elements $(Z_{it})_{i=1,2,\dots,N}$, for $t = 1, 2, \dots, T$. Then stack the vectors Z_t , for $t = 1, 2, \dots, T$, into an NT -vector Z , and similarly for ξ , $\tilde{\mu}$, and η (so that $\tilde{\mu} = \mathbf{1} \otimes \mu$, \otimes denoting the Kronecker product).

Conditional on ξ , the random vector $Z - \mu$ is multivariate normal with mean vector ξ and covariance matrix

$$G := FCov(\eta)F' = F\text{diag}(\Sigma_\eta, \dots, \Sigma_\eta)F' = F(I \otimes \Sigma_\eta)F',$$

where F is the $NT \times NT$ lower triangular block matrix with I_N (the $N \times N$ identity) on the diagonal, and whose t -th row is $[\Phi^t \ \Phi^{t-1} \ \dots \ \Phi \ I_N \ \mathbf{0} \ \dots \ \mathbf{0}]$, for $t = 0, 1, \dots, T - 1$. Thus the conditional density of Z , given ξ , is

$$f_{Z|\xi}(z) = \frac{1}{(2\pi \det G)^{NT/2}} e^{-\frac{1}{2}(z - \tilde{\mu} - \xi)' G^{-1} (z - \tilde{\mu} - \xi)}, \quad z \in \mathbb{R}^{NT}. \quad (3)$$

From this we can find the density of Z by multiplying with the density of ξ , then integrating out ξ .

To find the density of ξ , take arbitrary vectors x_1, \dots, x_T in \mathbb{R}^N , and write their elements as $x_t = (x_{it})_{i=1,2,\dots,N}$, for $t = 1, 2, \dots, T$. Since $\xi_t =$

$[\log \varepsilon_{1t}^2, \dots, \log \varepsilon_{Nt}^2]'$, and the ξ_t are i.i.d. over $t = 1, 2, \dots, T$, having distribution the same as ξ_1 , we have

$$\begin{aligned}
 F_{\xi}(\mathbf{x}; \Sigma_{\varepsilon}) &:= P(\xi_1 \leq \mathbf{x}_1, \dots, \xi_T \leq \mathbf{x}_T) = \prod_{t=1}^T P(\xi_1 \leq \mathbf{x}_t) \\
 &= \prod_{t=1}^T P(\log \varepsilon_{1t}^2 \leq x_{1t}, \dots, \log \varepsilon_{Nt}^2 \leq x_{Nt}) \\
 &= \prod_{t=1}^T P(|\varepsilon_{1t}| \leq e^{x_{1t}/2}, \dots, |\varepsilon_{Nt}| \leq e^{x_{Nt}/2}) \\
 &= \prod_{t=1}^T \int_{|u_1| \leq e^{x_{1t}/2}} \cdots \int_{|u_N| \leq e^{x_{Nt}/2}} \phi(\mathbf{u}; \Sigma_{\varepsilon}) d\mathbf{u}, \tag{4}
 \end{aligned}$$

where $\mathbf{u} = [u_1 \dots u_N] \in \mathbb{R}^N$, and $\phi(\cdot; \Sigma)$ is the density of an N -dimensional normally distributed random vector with mean vector $\mathbf{0}$ and covariance matrix Σ ; i.e.,

$$\phi(\mathbf{u}; \Sigma) = \frac{1}{(2\pi \det \Sigma)^{N/2}} e^{-\frac{1}{2} \mathbf{u}' \Sigma^{-1} \mathbf{u}}, \quad \mathbf{u} \in \mathbb{R}^N.$$

Expression (4) gives the joint cumulative distribution function of the elements in ξ . Write the integral in (4) as

$$\int_{-e^{x_{1t}/2}}^{e^{x_{1t}/2}} \cdots \int_{-e^{x_{Nt}/2}}^{e^{x_{Nt}/2}} \phi(\mathbf{u}; \Sigma_{\varepsilon}) d\mathbf{u} = \prod_{i=1}^N \left(\int_0^{e^{x_{it}/2}} + \int_{-e^{x_{it}/2}}^0 \right) \phi(\mathbf{u}; \Sigma_{\varepsilon}) d\mathbf{u},$$

and differentiate this with respect to x_{1t}, \dots, x_{Nt} , in turn, to get the density as

$$\frac{1}{2^N} e^{(x_{1t} + \dots + x_{Nt})/2} \sum^* \phi(\pm e^{x_{1t}/2}, \dots, \pm e^{x_{Nt}/2}; \Sigma_{\varepsilon}), \tag{5}$$

where \sum^* denotes summation over the 2^N possible ways of assigning + or - signs to the arguments of ϕ .

Substitute (5) in (4) to get the density of ξ as

$$f_{\xi}(\mathbf{x}; \Sigma_{\varepsilon}) = \frac{1}{2^{NT}} e^{\frac{1}{2} \sum_{t=1}^T \sum_{i=1}^N x_{it}} \prod_{t=1}^T \left(\sum^* \phi(\pm e^{x_{1t}/2}, \dots, \pm e^{x_{Nt}/2}; \Sigma_{\varepsilon}) \right). \tag{6}$$

Finally, multiplying (3) by (6) and integrating gives the density of \mathbf{Z} , that is, the likelihood of the data, as

$$\begin{aligned}
 f_{\mathbf{Z}}(\mathbf{z}) &= \frac{1}{2^{NT}} \frac{1}{(2\pi \det \mathbf{G})^{NT/2}} \int_{\mathbf{x}_1 \in \mathbb{R}^N} \cdots \int_{\mathbf{x}_T \in \mathbb{R}^N} e^{-\frac{1}{2} (\mathbf{z} - \tilde{\mu} - \xi)' \mathbf{G}^{-1} (\mathbf{z} - \tilde{\mu} - \xi)} e^{\frac{1}{2} \sum_{t=1}^T \sum_{i=1}^N x_{it}} \\
 &\times \prod_{t=1}^T \left(\sum^* \phi(\pm e^{x_{1t}/2}, \dots, \pm e^{x_{Nt}/2}; \Sigma_{\varepsilon}) \right) d\mathbf{x}_1 \cdots d\mathbf{x}_N, \quad \mathbf{z} \in \mathbb{R}^{NT}, \tag{7}
 \end{aligned}$$

where $\mathbf{x}_t = [x_{1t} \dots x_{Nt}]$ are arbitrary vectors in \mathbb{R}^N , $t = 1, 2, \dots, T$, and $\mathbf{x} = [\mathbf{x}'_1 \dots \mathbf{x}'_T]'$.

While expression (7) is quite explicit and in principle calculable for a given choice of the parameters, maximising it over the allowable range of the parameters is extremely difficult to do by standard means, and in fact intractable given current computational methods. The maximum of the function occurs, typically, in a very low probability region of the high dimensional space over which maximisation is taking place, and is hard to locate. A good exposition of the issues involved here is in Liesenfeld and Richard (2006). This is why we turn to the Bayesian and MCMC methods. As mentioned in Section 3, by taking uniform priors with large domains, the MCMC method produces estimates close to the MLEs, in any case.

Finally, we remark that the above analysis using normal distributions for the ε_t and η_t can be modified to allow for t -distributions. But, again, computationally intractable expressions result.

References

- Ahn D, Boudoukh J, Richardson M, Whitelaw R (2002) Partial adjustment or stale prices? Implications from stock index and futures return autocorrelations. *Rev Financ Stud* 15: 655–689
- Andersen T (1996) Return volatility and trading volume: an information flow interpretation of stochastic volatility. *J Finance* 51:169–204
- Broto C, Ruiz E (2004) Estimation methods for stochastic volatility models: a survey. *J Econ Surv* 18:613–649
- Chance D (1990) The effects of margins on the volatility of stock and derivative markets: a review of the evidence. New York University Monograph Series in Finance and Economics II
- Chib S, Nardari F, Shephard N (2002) Markov chain Monte Carlo methods for stochastic volatility models. *J Econom* 108:281–316
- Darrat A, Zhong M (2005) Equity market linkage and multinational trade accords: the case of NAFTA. *J Int Money Financ* 24:793–817
- ESTIMA (2004) RATS user's guide (version 6). ESTIMA, Evanston
- Ewing B, Payne J, Sowell C (1999) NAFTA and North American stock market linkages: an empirical note. *N Am J Econ Financ* 10:443–451
- Fleming J, Kirby C, Ostdiek B (1998) Information and volatility linkages in the stock, bond and money markets. *J Financ Econ* 49:111–137
- Gelman A, Carlin JB, Stern HS, Rubin DB (1995) Bayesian data analysis. Chapman and Hall, London
- Gelman A, Meng X-L, Stern H (1996) Posterior predictive assessment of model fitness via realized discrepancies. *Stat Sin* 6:733–807
- Gilks WR, Richardson S, Spiegelhalter DJ (1996) Markov chain Monte Carlo in practice. Chapman and Hall, London
- Gould DM (1996) Distinguishing NAFTA from the peso crisis. *Southwest Economy*, Federal Reserve Bank of Dallas 6–10
- Gould DM (1998) Has NAFTA changed North American trade? *Economic Review*, Federal Reserve Bank of Dallas 12–23
- Grossman SJ, Stiglitz J (1980) On the impossibility of informationally efficient markets. *Am Econ Rev* 70:393–408
- Harvey A, Ruiz E, Shephard N (1994) Multivariate stochastic variance models. *Rev Econ Stud* 61:247–264
- Hufbauer GC, Schott JJ (1993) NAFTA: an assessment. Peterson Institute for International Economics, Washington DC

- Hufbauer GC, Schott JJ (2005) NAFTA Revisited: Achievements and Challenges. Peterson Institute for International Economics, Washington DC
- Jacquier E, Polson NG, Rossi PE (1994) Bayesian analysis of stochastic volatility models. *J Econ Bus Stat* 12:371–389
- Kodres L, Pritsker M (2002) A rational expectations model of financial contagion. *J Finance* 57:769–799
- Kose M, Meredith G, Towe C (2005) How has NAFTA affected the Mexican economy? Review and evidence. In: Langhammer RJ, Vinhas de Souza L (eds) *Monetary policy and macroeconomic stabilization in Latin America*. Springer, Berlin Heidelberg New York, pp 35–81
- Lamoureux C, Lastrapes W (1990) Heteroskedasticity in stock return data: volume versus GARCH effects. *J Finance* 45:221–229
- Lederman D, Menendez A, Perry G, Stiglitz J (2003) Mexican investment after the tequila crisis: basic economics, “confidence” effects or market imperfections? *J Int Money Financ* 22: 131–151
- Liesenfeld R, Richard J-F (2006) Classical and Bayesian analysis of univariate and multivariate stochastic volatility models. *Econom Rev* 25:335–360
- Lunn DJ, Thomas A, Best N, Spiegelhalter D (2000) WinBUGS—A Bayesian modelling framework: concepts, structure and extensibility. *Stat Comput* 10:325–337
- Robert CP, Casella G (2000) *Monte Carlo statistical methods*. Springer, New York
- Ruiz E (1994) Quasi-Maximum likelihood estimation of stochastic variance models. *J Econom* 63:284–306
- Spiegelhalter D, Thomas A, Best N, Gilks W (2003) WinBUGS user manual (version 1.4). MRC Biostatistics Unit, Cambridge, UK
- Stiglitz J (2004) The broken promise of NAFTA. *New York Times* 6, January
- Stoll H, Whaley R (1990) The dynamics of stock index and stock index futures returns. *J Financ Quant Anal* 25:441–468
- Tauchen G, Pitts M (1983) The price variability-volume relationship on speculative markets. *Econometrica* 51:485–505
- Taylor SJ (1986) *Modelling financial time series*. Wiley, Chichester
- Taylor SJ (1994) Modelling stochastic volatility: a review and comparative study. *Math Financ* 4:183–204
- Wongswan J (2006) Transmission of information across international equity markets. *Rev Financ Stud* 19:1157–1189