TABLE IV
10 KEY WORLD EVENTS AFFECTING THE IBOVESPA STOCK INDEX (SAO PAULO STOCK EXCHANGE) OVER THE PERIOD OF 01/03/1997 TO 01/16/2001, AS CITED BY CARVALHO AND LOPES [7]

Date	Event
07/02/1997	Thailand devalues the Baht by as much as 20%
08/11/1997	IMF and Thailand set a rescue agreement
10/23/1997	Hong Kongs stock index falls 10.4%. South Korea won starts to weaken
12/02/1997	IMF and South Korea set a bailout agreement
06/01/1998	Russias stock market crashes
06/20/1998	IMF gives final approval to a loan package to Russia
08/19/1998	Russia officially falls into default
10/09/1998	IMF and World Bank joint meeting to discuss global economic crisis. The Fed cuts interest rates
01/15/1999	The Brazilian government allows its currency, the Real, to float freely by lifting exchange controls
02/02/1999	Arminio Fraga is named President of Brazils Central Bank

volatility (MSSV) model. In the maneuvering target tracking community, it is often further assumed that the dynamic matrix is shared and *known* (due to the understood physics of the target). We explore both of these variants in the following sections.

A. Shared Dynamic Matrix, Switching Driving Noise

In many applications, the dynamics of the switching process can be described by a shared linear dynamical system matrix A; the dynamics within a given mode are then determined by some external force acting upon this LDS and it is how this force is exerted that is mode-specific. The general form for such an SLDS is given by

$$\begin{aligned} z_t \mid z_{t-1} &\sim \pi_{z_{t-1}} \\ \boldsymbol{x}_t &= A\boldsymbol{x}_{t-1} + \boldsymbol{e}_t^{(z_t)} \\ \boldsymbol{y}_t &= C\boldsymbol{x}_t + \boldsymbol{w}_t \end{aligned} \tag{37}$$

with process and measurement noise $\boldsymbol{e}_t^{(k)} \sim \mathcal{N}\left(\boldsymbol{\mu}^{(k)}, \Sigma^{(k)}\right)$ and $\boldsymbol{w}_t \sim \mathcal{N}(0, R)$, respectively. In this scenario, the data are generated from one dynamic matrix, A and multiple process noise covariance matrices, $\Sigma^{(k)}$. Thus, one cannot place a MNIW prior jointly on these parameters (conditioned on $\boldsymbol{\mu}^{(k)}$) due to the coupling of the parameters in this prior. We instead consider independent priors on A, $\Sigma^{(k)}$ and $\boldsymbol{\mu}^{(k)}$. We will refer to the choice of a normal prior on A, inverse-Wishart prior on $\Sigma^{(k)}$ and normal prior on $\boldsymbol{\mu}^{(k)}$ as the *N-IW-N* prior. See [42] for details on deriving the resulting posterior distributions given these independent priors.

1) Stochastic Volatility: An example of an SLDS in a similar form to that of (37) is the Markov switching stochastic volatility (MSSV) model [5], [6], [44]. The MSSV assumes that the log-volatilities follow an AR(1) process with a Markov switching mean. This underlying process is observed via conditionally independent and normally distributed daily returns. Specifically, let y_t represent, for example, the daily returns of a stock index. The state x_t is then given the interpretation of log-volatilities and the resulting state space model is given by [7]

$$\begin{aligned} z_t \mid z_{t-1} &\sim \pi_{z_{t-1}} \\ x_t &= a x_{t-1} + e_t^{(z_t)} \\ y_t &= u_t(x_t) \end{aligned} \tag{38}$$

with $e_t^{(k)} \sim \mathcal{N}\left(\mu^{(k)}, \sigma^2\right)$ and $u_t(x_t) \sim \mathcal{N}(0, \exp(x_t))$. Here, only the mean of the process noise is mode-specific. Note, however, that the measurement equation is non-linear in the state x_t. Carvalho and Lopes [7] employ a particle filtering approach to cope with these non-linearities. In [6], the MSSV is instead modeled in the log-squared-daily-returns domain such that

$$\log(y_t^2) = x_t + w_t \tag{39}$$

where w_t is additive, non-Gaussian noise. This noise is sometimes approximated by a moment-matched Gaussian [45], while So *et al.* [6] use a mixture of Gaussians approximation. The MSSV is then typically bestowed a fixed set of two or three regimes of volatility.

We examine the IBOVESPA stock index (Sao Paulo Stock Exchange) over the period of 01/03/1997 to 01/16/2001, during which ten key world events are cited in [7] as affecting the emerging Brazilian market. The key world events are summarized in Table IV and shown in the plots of Fig. 8. Use of this dataset was motivated by the work of Carvalho and Lopes [7], in which a two-mode MSSV model is assumed. We consider a variant of the HDP-SLDS to match the MSSV model of (38). Specifically, we examine log-squared daily returns, as in (39) and use a DP mixture of Gaussians to model the measurement noise

$$\begin{aligned} e_t^{(k)} &\sim \mathcal{N}\left(\mu^{(k)}, \Sigma^{(k)}\right) \\ w_t &\sim \sum_{\ell=1}^{\infty} \omega_\ell \mathcal{N}(0, R_\ell) \\ \boldsymbol{\omega} &\sim \mathrm{GEM}(\sigma_r) \\ R_\ell &\sim \mathrm{IW}(n_r, S_r). \end{aligned} \tag{40}$$

We truncate the measurement noise DP mixture to 10 components. For the HDP concentration hyperparameters, α, γ and κ, we use the same prior distributions as in Sections IV-A–IV-C. For the dynamic parameters, we rely on the N-IW-N prior described in Section V-A and once again set the hyperparameters of this prior from statistics of the data as described in the Appendix. Since we allow for a mean on the process noise and examine log-squared daily returns, we do not preprocess the data.

The posterior probability of an HDP-SLDS inferred change point is shown in Fig. 8(a) and in Fig. 8(b) we display the corresponding plot for a non-sticky variant (i.e., with $\kappa = 0$ so that there is no bias towards mode self-transitions.) The HDP-SLDS is able to infer very similar change points to those presented in [7]. Without the sticky extension, the nonsticky model variant oversegments the data and rapidly switches between redundant states leading to many inferred change points that do not align with any world event. As a quantitative comparison of the inferred change points, we compute a Hamming distance metric