

**Likelihood evaluation of jump-diffusion models using deterministic nonlinear filters**

J.-F. Bégin,  
M. Boudreault

G-2019-38

June 2019

---

La collection *Les Cahiers du GERAD* est constituée des travaux de recherche menés par nos membres. La plupart de ces documents de travail a été soumis à des revues avec comité de révision. Lorsqu'un document est accepté et publié, le pdf original est retiré si c'est nécessaire et un lien vers l'article publié est ajouté.

**Citation suggérée :** J.-F. Bégin, M. Boudreault (Juin 2019). Likelihood evaluation of jump-diffusion models using deterministic nonlinear filters, Rapport technique, Les Cahiers du GERAD G-2019-38, GERAD, HEC Montréal, Canada.

**Avant de citer ce rapport technique**, veuillez visiter notre site Web (<https://www.gerad.ca/fr/papers/G-2019-38>) afin de mettre à jour vos données de référence, s'il a été publié dans une revue scientifique.

---

La publication de ces rapports de recherche est rendue possible grâce au soutien de HEC Montréal, Polytechnique Montréal, Université McGill, Université du Québec à Montréal, ainsi que du Fonds de recherche du Québec – Nature et technologies.

Dépôt légal – Bibliothèque et Archives nationales du Québec, 2019  
– Bibliothèque et Archives Canada, 2019

The series *Les Cahiers du GERAD* consists of working papers carried out by our members. Most of these pre-prints have been submitted to peer-reviewed journals. When accepted and published, if necessary, the original pdf is removed and a link to the published article is added.

**Suggested citation:** J.-F. Bégin, M. Boudreault (June 2019). Likelihood evaluation of jump-diffusion models using deterministic nonlinear filters, Technical report, Les Cahiers du GERAD G-2019-38, GERAD, HEC Montréal, Canada.

**Before citing this technical report**, please visit our website (<https://www.gerad.ca/en/papers/G-2019-38>) to update your reference data, if it has been published in a scientific journal.

---

The publication of these research reports is made possible thanks to the support of HEC Montréal, Polytechnique Montréal, McGill University, Université du Québec à Montréal, as well as the Fonds de recherche du Québec – Nature et technologies.

Legal deposit – Bibliothèque et Archives nationales du Québec, 2019  
– Library and Archives Canada, 2019



# Likelihood evaluation of jump-diffusion models using deterministic nonlinear filters

Jean-François Bégin <sup>a</sup>

Mathieu Boudreault <sup>b,c</sup>

<sup>a</sup> Department of Statistics and Actuarial Science, Simon Fraser University, Burnaby (British Columbia), Canada, V5A 1S6

<sup>b</sup> GERAD, Montréal (Québec), Canada, H3T 2A7

<sup>c</sup> Department of Mathematics, Université du Québec à Montréal, Montréal (Québec), Canada, H2X 3Y7

[jbegin@sfu.ca](mailto:jbegin@sfu.ca)

[boudreault.mathieu@uqam.ca](mailto:boudreault.mathieu@uqam.ca)

June 2019

Les Cahiers du GERAD

G–2019–38

Copyright © 2019 GERAD, Bégin, Boudreault

---

Les textes publiés dans la série des rapports de recherche *Les Cahiers du GERAD* n'engagent que la responsabilité de leurs auteurs. Les auteurs conservent leur droit d'auteur et leurs droits moraux sur leurs publications et les utilisateurs s'engagent à reconnaître et respecter les exigences légales associées à ces droits. Ainsi, les utilisateurs:

- Peuvent télécharger et imprimer une copie de toute publication du portail public aux fins d'étude ou de recherche privée;
- Ne peuvent pas distribuer le matériel ou l'utiliser pour une activité à but lucratif ou pour un gain commercial;
- Peuvent distribuer gratuitement l'URL identifiant la publication.

Si vous pensez que ce document enfreint le droit d'auteur, contactez-nous en fournissant des détails. Nous supprimerons immédiatement l'accès au travail et enquêterons sur votre demande.

The authors are exclusively responsible for the content of their research papers published in the series *Les Cahiers du GERAD*. Copyright and moral rights for the publications are retained by the authors and the users must commit themselves to recognize and abide the legal requirements associated with these rights. Thus, users:

- May download and print one copy of any publication from the public portal for the purpose of private study or research;
- May not further distribute the material or use it for any profit-making activity or commercial gain;
- May freely distribute the URL identifying the publication.

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

**Abstract:** In this study, we develop a deterministic nonlinear filtering algorithm based on a high-dimensional version of Kitagawa (1987) to evaluate the likelihood function of models that allow for stochastic volatility and jumps whose arrival intensity is also stochastic. We show numerically that the deterministic filtering method is precise and much faster than the particle filter, in addition to yielding a smooth function over the parameter space. We then find the maximum likelihood estimates of various models that include stochastic volatility, jumps in the returns and variance, and also stochastic jump arrival intensity with the S&P 500 daily returns. During the Great Recession, the jump arrival intensity increases significantly and contributes to the clustering of volatility and negative returns.

**Keywords:** Discrete nonlinear filtering, particle filter, sequential importance resampling, stochastic volatility, maximum likelihood estimation

---

**Acknowledgments:** For their useful comments, the authors would like to thank Maciej Augustyniak, Geneviève Gauthier, Frédéric Godin, as well as the seminar participants at the 2019 Quantact Workshop on Financial Mathematics and the 23rd International Congress on Insurance: Mathematics and Economics. Bégin wishes to acknowledge the financial support of the National Science and Engineering Research Council of Canada (NSERC), Simon Fraser University and the NVIDIA Corporation. Boudreault would like to thank the financial support of NSERC.

# 1 Introduction and motivation

Over the last thirty years, jump-diffusion models (with stochastic volatility) have become increasingly popular for their ability to replicate important stylized facts such as heavy tails, no autocorrelation in the returns, volatility clustering, etc. (Cont, 2007). Yet, parameter estimation of these models is cumbersome as stochastic volatility and jumps are latent—or not directly observed. Therefore, the first studies that looked into that problem focused mostly on stochastic volatility models using the (generalized) method of moments, quasi-, simulated and approximate maximum likelihood, and Bayesian methods.<sup>1</sup>

Most recent progress in parameter estimation of jump-diffusion models has been accomplished with numerical recursive prediction-update algorithms, namely sequential Monte Carlo (SMC) methods and discrete nonlinear filters (DNF). Such methods provide the posterior distribution of latent state variables conditional on current and past observations in a recursive manner.<sup>2</sup> As a by-product of the method, one can obtain the likelihood function. On the one hand, sequential Monte Carlo (SMC) methods—so-called particle filters—as introduced by Gordon et al. (1993), have been used to compute and maximize the likelihood function (Johannes et al., 2009; Christoffersen et al., 2010; Pitt et al., 2014; Bégin et al., 2019; Bardgett et al., 2019). Although the SMC method is very flexible, it is Monte Carlo-based and thus computationally intensive. Moreover, the likelihood function is random and not smooth for a finite set of particles, which can be a significant issue for frequentist-based inference.<sup>3</sup>

On the other hand, one can approximate the prediction density as well as the likelihood function with numerical integration schemes; this is the framework put forward by Kitagawa (1987). As such, the dimension of the integration problem increases with the number of state variables which explains why many authors solely focused on stochastic volatility models (Fridman and Harris, 1998; Watanabe, 1999; Bartolucci and De Luca, 2001; Clements et al., 2006a; Langrock et al., 2012) whereas jumps were ignored.

In this article, we present a DNF-based method for the likelihood evaluation of stochastic volatility models that include return jumps, variance jumps, as well as stochastic jump arrival intensity. Those include well-known models such as the stochastic volatility (SV) model, stochastic volatility with return jumps (SVYJ) model, and the stochastic volatility with return and variance correlated jumps (SVCJ) model, which are nested cases of the general model. We then assess the accuracy of the proposed methodology and compare it with the most common SMC method—the sequential importance resampling (SIR) of Gordon et al. (1993). Finally, we apply the method to a typical financial time series: the S&P 500 Index.

Our main contributions and findings are as follows. We first develop a prediction-update algorithm for the SVCJSI model—and its nested models. The proposed approach reduces the dimensionality of the resulting integrals when compared to a naive application of Kitagawa (1987). As a result, the proposed DNF method is both accurate and faster than the SIR. And because the DNF yields a smooth likelihood function in the parameter space, it is ideal for numerical maximization. We thus compute the maximum likelihood estimates of the SV, SVYJ, SVCJ and SVCJSI<sup>4</sup> models using S&P 500 daily returns. We find that during the Great Recession, the jump arrival intensity increases significantly and contributes to the clustering of volatility and negative returns.

---

<sup>1</sup>The first study looking into parameter estimation of stochastic volatility models is Taylor (1986) which was based upon the method of moments. Generalizations of the method of moments were proposed by Melino and Turnbull (1990), Andersen and Sørensen (1996), Pan (2002). Quasi-maximum likelihood approaches were suggested by Nelson (1988) and Harvey et al. (1994); Danielsson (1994) and Brandt and Santa-Clara (2002) used simulated maximum likelihood whereas approximate maximum likelihood were introduced by Aït-Sahalia (2002), Bates (2006) and Aït-Sahalia and Kimmel (2007). Finally, another family of approaches based upon the Bayesian paradigm was used in this literature, namely Shephard (1993), Jacquier et al. (1994), Johannes et al. (1999) and Eraker (2001).

<sup>2</sup>A classic example of prediction-update algorithms is the Kalman (1960) filter which provides the exact posterior distribution of the state variables in a Gaussian and linear framework.

<sup>3</sup>Malik and Pitt (2011) introduced a resampling algorithm that smoothes the likelihood function but according to Creal (2012), this only works when the state dimension is one, which is typically the case when only stochastic volatility is included whereas jumps are excluded.

<sup>4</sup>SVCJ with stochastic jump arrival intensity

The article is organized as follows. Section 2 presents the econometric framework. The DNF is explained in Section 3. Section 4 discusses the accuracy and the reliability of the DNF method. The speed-accuracy trade-off is assessed in Section 5. An empirical application is presented in Section 6. Finally, some concluding remarks are provided in Section 7.

## 2 Framework

This section describes the continuous- and discrete-time frameworks. We then specifically introduce the four models that are used throughout this paper, namely the SV, SVYJ, SVCJ, and the full model with stochastic jump arrival intensity, known as SVCJSI.

### 2.1 Continuous-time framework

We fix a filtered probability space  $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$  and a filtration  $\mathbb{F} = \{\mathcal{F}_t : t \geq 0\}$  satisfying the usual conditions. Let  $S_t$  be the time- $t$  price of a security,  $V_t$  the time- $t$  instantaneous variance and  $\Lambda_t$  the time- $t$  jump intensity. Under the objective measure  $\mathbb{P}$ , their dynamics are given by the following equations:

$$\begin{aligned} \frac{dS_t}{S_{t-}} &= (\mu - \bar{\alpha}\Lambda_t) dt + \sqrt{V_{t-}} dW_t^S + d\left(\sum_{n=1}^{N_t} (e^{Z_n^S} - 1)\right), \\ dV_t &= \kappa(\theta - V_{t-}) dt + \sigma\sqrt{V_{t-}} dW_t^V + d\left(\sum_{n=1}^{N_t} Z_n^V\right), \\ d\Lambda_t &= \chi(\omega - \Lambda_t) dt + \xi\sqrt{\Lambda_t} dW_t^\Lambda, \end{aligned}$$

where  $W^S$ ,  $W^V$  and  $W^\Lambda$  are three Brownian motions such that  $d\langle W^S, W^V \rangle_t = \rho_v dt$ ,  $d\langle W^S, W^\Lambda \rangle_t = \rho_\lambda dt$  and  $d\langle W^V, W^\Lambda \rangle_t = 0 dt$ . Moreover,  $\{N_t\}_{t \geq 0}$  is a Cox process (doubly stochastic Poisson process) with jump arrival intensity  $\Lambda_t$  and jump sizes to be further discussed. Note that only  $S_t$  is observed whereas  $V_t$  and  $\Lambda_t$  are unobserved or latent.

The first stochastic differential equation (SDE) resembles that proposed by Merton (1976) as it includes both a jump process as well as a diffusive component. Yet, in opposition to Merton, our diffusive component allows for stochastic volatility—modelled by the second SDE. The stochastic variance SDE is given by the square-root process—similar to that used by Heston (1993)—to which variance jumps are added. Parameter  $\kappa$  denotes the speed of mean reversion,  $\theta$  is the unconditional mean variance, and  $\sigma$  is the so-called variance of the variance parameter. Finally, the last SDE models the stochastic jump arrival intensity; parameter  $\chi$  is related to the speed of mean reversion,  $\omega$  is the long run level of the stochastic jump intensity and  $\xi$  is the variance of the intensity parameter.

We assume that each of the return jumps  $\{Z_n^S\}_{n \in \mathbb{N}}$  are normally distributed with a mean of  $\alpha + \rho_z Z_n^V$  and a standard deviation of  $\delta$ , i.e.,  $Z_n^S \sim \mathcal{N}(\alpha + \rho_z Z_n^V, \delta^2)$  and that the variance jumps  $\{Z_n^V\}_{n \in \mathbb{N}}$  are exponentially distributed with a mean of  $\nu$ , i.e.,  $Z_n^V \sim \text{Exp}(\nu)$ . Similar return jump innovations have been used by Bates (1996), Bakshi et al. (1997), Duffie et al. (2000), Pan (2002), Eraker et al. (2003), and Johannes et al. (2009) in continuous-time models whereas Maheu and McCurdy (2004), Christoffersen et al. (2012) and Ornathanalai (2014) (among others) have introduced normally distributed jumps in discrete-time models. Evidence for the presence of positive jumps in the variance process has been found by numerous authors, namely Bates (2000), Duffie et al. (2000), Pan (2002), Eraker et al. (2003) and Todorov and Tauchen (2011).

Finally, the term  $\Lambda_t \bar{\alpha}$  is the jump compensator that makes sure that  $\mathbb{E}^\mathbb{P}[S_T | \mathcal{F}_t] = S_t e^{\mu(T-t)}$ , with  $\mu$  being the annual expected rate of return on the stock in our setting. We thus have  $\bar{\alpha} = \exp(\alpha + \frac{1}{2}\delta^2)/(1 - \rho_z \nu) - 1$ .

## 2.2 Discrete-time framework

As we typically observe most continuous-time processes only in discrete-time, we thus need to discretize our continuous-time dynamics. Setting  $h$  as the length of a time interval (say e.g. a month, a week or a calendar/trading day), then we let  $y_t \equiv \log(S_{th}/S_{(t-1)h})$  be the time- $t$  (continuously compounded ex-dividend) return of a security. As such, the information set available to the econometrician  $\mathbb{G} = \{\mathcal{G}_t : t \in \mathbb{N}\}$  is coarser than that of the model, i.e.,  $\mathcal{G}_t \subseteq \mathcal{F}_t$  since  $\mathcal{G}_t = \sigma(\{y_s\}_{s=1}^t)$ . Similarly, we define  $v_t$  and  $\lambda_t$  as the time- $th$  discretized variance and jump intensity processes, respectively. Again, as in the continuous-time version of the model,  $v_t$  and  $\lambda_t$  are latent.

The Euler-Maruyama discretization scheme applied to the continuous-time version of the model yields

$$\begin{aligned} y_t &= \left( \mu - \frac{1}{2}v_{t-1} - \bar{\alpha}\lambda_{t-1} \right) h + \sqrt{v_{t-1}h} \varepsilon_t^y + \sum_{i=1}^{n_t} z_{t,i}^y, \\ v_t &= v_{t-1} + \kappa(\theta - v_{t-1})h + \sigma\sqrt{v_{t-1}h} \varepsilon_t^v + \sum_{i=1}^{n_t} z_{t,i}^v, \\ \lambda_t &= \lambda_{t-1} + \chi(\omega - \lambda_{t-1})h + \xi\sqrt{\lambda_{t-1}h} \varepsilon_t^\lambda, \end{aligned}$$

where  $n_t = N_{th} - N_{(t-1)h}$  is given by a Poisson random variable with parameter  $\lambda_{t-1}h$ . Moreover, the standardized innovations are normally distributed, i.e.,  $\varepsilon_t^i \sim \mathcal{N}(0, 1)$  for  $i \in \{\perp, v, \lambda\}$  with

$$\varepsilon_t^y = \sqrt{1 - \rho_v^2 - \rho_\lambda^2} \varepsilon_t^\perp + \rho_v \varepsilon_t^v + \rho_\lambda \varepsilon_t^\lambda$$

obtained from the Cholesky decomposition. Finally, return jumps are also normally distributed, i.e.  $z_{t,i}^y \sim \mathcal{N}(\alpha + \rho z_{t,i}^v, \delta^2)$ , whereas the distribution of variance jumps is still exponential, i.e.,  $z_{t,i}^v \sim \text{Exp}(\nu)$ .

## 2.3 Nested models

This general specification embeds numerous stochastic volatility-type models. The four specific models considered in this study are detailed below:

- **SV**: Stochastic volatility model with no jumps, obtained by letting  $\chi = \omega = \xi = 0$ . It is similar to the model proposed by Heston (1993) in the sense that it incorporates a leverage effect. Specifications alike were also investigated by Melino and Turnbull (1990), Danielsson and Richard (1993), Shephard (1993), and Jacquier et al. (1994), among others.
- **SVYJ**: Stochastic volatility model with jumps in return only, obtained by letting  $\chi = 1$  and  $\xi = \nu = 0$ . It resembles the Bates's (1996) model. In econometrics, Pitt et al. (2014) proposed the use of jumps in the return's dynamics.
- **SV CJ**: Stochastic volatility with simultaneous and correlated jumps in return and variance, obtained by letting  $\chi = 1$  and  $\xi = 0$ . The specification was first used by Duffie et al. (2000).
- **SV CJ SI**: Stochastic volatility with simultaneous and correlated jumps in return and variance along with stochastic jump arrival intensity. This is the most complicated and general specification used in this study.

## 3 Discrete nonlinear filtering

Because the volatility, jumps and the jump intensity factors are latent variables in the model, it is difficult to evaluate the likelihood function  $\mathcal{L}$  in a direct manner. As such, we propose in this section a deterministic DNF scheme that uses a recursive prediction-update algorithm in the spirit of Kitagawa (1987). This methodology is based on a discretization of the state-space continuous latent variables, similar to the method proposed by Fridman and Harris (1998), Bartolucci and De Luca (2001), Clements et al. (2006b), and Langrock et al. (2012). It is important to note that such a

methodology has not been applied to the likelihood evaluation of SVYJ, SVCJ and SVCJSI models, which is a key contribution of this paper.

Let  $\mathbf{x}_t \in \mathcal{X}$  be the latent state variables at time  $t$  where  $\mathbf{x}_t$  is either a scalar or a vector. Moreover, let us define by  $y_t, t = 1, 2, \dots, T$  the time- $t$  observations—i.e., returns in this study. Series of latent factors and observations are denoted by  $\mathbf{x}_{0:t} \equiv \{\mathbf{x}_s\}_{s=0}^t$  and  $y_{1:t} \equiv \{y_s\}_{s=1}^t$ . For a parameter set  $\Theta$ , the likelihood function of  $\Theta$  is

$$\mathcal{L}(\Theta) = f(y_{1:T}) = f(y_1) \prod_{t=2}^T f(y_t | y_{1:t-1}) = \int_{\mathcal{X}^{T+1}} f(\mathbf{x}_{0:T}, y_{1:T}) d\mathbf{x}_{0:T}, \quad (1)$$

which involves a multidimensional integration problem.

### 3.1 General methodology

The one-step ahead *prediction* of the latent states' distribution (conditional on the past returns) is given by

$$f(\mathbf{x}_t | y_{1:t-1}) = \int_{\mathcal{X}} q(\mathbf{x}_t | \mathbf{x}_{t-1}, y_{1:t-1}) u(\mathbf{x}_{t-1} | y_{1:t-1}) d\mathbf{x}_{t-1},$$

where  $q$  is the transition probability distribution of  $\mathbf{x}_t$  given  $\mathbf{x}_{t-1}$  and the past returns. Once a new return is available, the *posterior* distribution of the state variables at time  $t$ , conditional on  $\mathcal{G}_t$ , may now be obtained as

$$u(\mathbf{x}_t | y_{1:t}) = \frac{r(y_t | \mathbf{x}_t, y_{1:t-1}) f(\mathbf{x}_t | y_{1:t-1})}{f(y_t | y_{1:t-1})}, \quad (2)$$

where  $r$  is the measurement probability density—or the conditional time- $t$  return density given the state space variables  $\mathbf{x}_t$ . This is also known as the *update* step of the algorithm. The denominator of Equation (2) is the time- $t$  likelihood contribution of return  $y_t$  conditional on the past returns, i.e.,

$$f(y_t | y_{1:t-1}) = \int_{\mathcal{X}} r(y_t | \mathbf{x}_t, y_{1:t-1}) f(\mathbf{x}_t | y_{1:t-1}) d\mathbf{x}_t.$$

Given the model formulation of Section 2,  $\mathbf{x}_t$  depends upon  $\mathbf{x}_{t-1}$ . Hence, the time- $t$  likelihood contribution can be rewritten as

$$\begin{aligned} f(y_t | y_{1:t-1}) &= \iint_{\mathcal{X}^2} r(y_t | \mathbf{x}_t, \mathbf{x}_{t-1}, y_{1:t-1}) f(\mathbf{x}_t, \mathbf{x}_{t-1} | y_{1:t-1}) d\mathbf{x}_t d\mathbf{x}_{t-1} \\ &= \iint_{\mathcal{X}^2} r(y_t | \mathbf{x}_t, \mathbf{x}_{t-1}, y_{1:t-1}) q(\mathbf{x}_t | \mathbf{x}_{t-1}, y_{1:t-1}) u(\mathbf{x}_{t-1} | y_{1:t-1}) d\mathbf{x}_t d\mathbf{x}_{t-1}, \end{aligned} \quad (3)$$

where  $u(\mathbf{x}_{t-1} | y_{1:t-1})$  is similar to Equation (2). Then, the posterior density of Equation (2) can also be obtained recursively using the same idea:

$$\begin{aligned} u(\mathbf{x}_t | y_{1:t}) &= \frac{\int_{\mathcal{X}} r(y_t | \mathbf{x}_t, \mathbf{x}_{t-1}, y_{1:t-1}) f(\mathbf{x}_t, \mathbf{x}_{t-1} | y_{1:t-1}) d\mathbf{x}_{t-1}}{f(y_t | y_{1:t-1})} \\ &= \frac{\int_{\mathcal{X}} r(y_t | \mathbf{x}_t, \mathbf{x}_{t-1}, y_{1:t-1}) q(\mathbf{x}_t | \mathbf{x}_{t-1}, y_{1:t-1}) u(\mathbf{x}_{t-1} | y_{1:t-1}) d\mathbf{x}_{t-1}}{f(y_t | y_{1:t-1})}. \end{aligned} \quad (4)$$

Therefore, to obtain the likelihood function of Equation (1), one needs to find the posterior distribution of the latent states and the likelihood contribution at all times. Algorithm 1 summarizes the various steps of the general DNF.

**Algorithm 1** Discrete nonlinear filter

- 
- 1: set  $u(\mathbf{x}_0) = p(\mathbf{x}_0)$ , where  $p$  is the initial state density
  - 2: **for** each  $t \in \{1, \dots, T\}$  **do**
  - 3:   get the likelihood contribution  $f(y_t | y_{1:t-1})$  using numerical integration
  - 4:   get the posterior distribution  $u(\mathbf{x}_t | y_{1:t})$  using numerical integration
  - 5: **end for**
  - 6: compute the likelihood function  $\mathcal{L}(\Theta)$  by taking the product of the likelihood contributions
- 

**3.2 Likelihood function of the SVCJSI model**

We seek to derive the likelihood function of the SVCJSI model using the DNF method.<sup>5</sup> Specifically, for the framework of Section 2, we have that  $\mathbf{x}_t = [v_t \ \lambda_t \ j_t^y \ j_t^v]$ , where

$$j_t^y = \sum_{i=1}^{n_t} z_{t,i}^y \quad \text{and} \quad j_t^v = \sum_{i=1}^{n_t} z_{t,i}^v,$$

as knowing every individual jump sizes over a day is futile—only the aggregated jump is important in the discretization. Therefore,  $\mathcal{X} = \mathbb{R}_+ \times \mathbb{R}_+ \times \mathbb{R} \times \mathbb{R}_+$  and the latent state vector is four-dimensional for the SVCJSI model.

Naively computing Equations (3) and (4) might be cumbersome given the high-dimensionality of the state variables. Fortunately, some simplifications can be applied to reduce the dimension of the problem. First, by conditioning on the number of jumps at time  $t$ , we can simplify the transition density in the following way:

$$\begin{aligned} q(\mathbf{x}_t | \mathbf{x}_{t-1}, y_{1:t-1}, n_t = n) &= q(v_t, \lambda_t, j_t^y, j_t^v | v_{t-1}, \lambda_{t-1}, j_{t-1}^y, j_{t-1}^v, y_{1:t-1}, n_t = n) \\ &= q(v_t | v_{t-1}, j_t^v) q(\lambda_t | \lambda_{t-1}) q(j_t^y | j_t^v, n_t = n) q(j_t^v | n_t = n) \end{aligned}$$

as  $v_t$  and  $\lambda_t$  depend upon  $v_{t-1}$  and  $\lambda_{t-1}$ , respectively. Yet,  $j_t^{(y)}$  and  $j_t^{(v)}$  do not depend on the previous value of the latent state; in fact, return and variance jump sizes are (iid) transient latent states. By conditioning again on the number of jumps at time  $t$ , the observation density is given by:

$$r(y_t | \mathbf{x}_t, \mathbf{x}_{t-1}, y_{1:t-1}, n_t = n) = r(y_t | v_t, v_{t-1}, \lambda_t, \lambda_{t-1}, j_t^y, j_t^v, n_t = n).$$

Then, following the rationale outlined at the beginning of Section 3, the time- $t$  likelihood contribution of Equation (3) can be rewritten as

$$\begin{aligned} f(y_t | y_{1:t-1}) &= \iint_{\mathcal{X}^2} r(y_t | \mathbf{x}_t, \mathbf{x}_{t-1}, y_{1:t-1}) q(\mathbf{x}_t | \mathbf{x}_{t-1}, y_{1:t-1}) u(\mathbf{x}_{t-1} | y_{1:t-1}) d\mathbf{x}_t d\mathbf{x}_{t-1} \\ &= \sum_{n=0}^{\infty} \iint_{\mathcal{X}^2} r(y_t | \mathbf{x}_t, \mathbf{x}_{t-1}, y_{1:t-1}, n_t = n) q(\mathbf{x}_t | \mathbf{x}_{t-1}, y_{1:t-1}, n_t = n) \\ &\quad u(\mathbf{x}_{t-1} | y_{1:t-1}) \mathbb{P}(n_t = n | \mathbf{x}_t, \mathbf{x}_{t-1}) d\mathbf{x}_t d\mathbf{x}_{t-1} \\ &= \sum_{n=0}^{\infty} \int \cdots \int_{\mathcal{X} \times \mathbb{R}_+ \times \mathbb{R}_+} \mathbb{P}(n_t = n | \lambda_{t-1}) r(y_t | v_t, v_{t-1}, \lambda_t, \lambda_{t-1}, j_t^y, j_t^v, n_t = n) \\ &\quad q(j_t^y | j_t^v, n_t = n) q(v_t | v_{t-1}, j_t^v) q(j_t^v | n_t = n) q(\lambda_t | \lambda_{t-1}) \\ &\quad u(v_{t-1}, \lambda_{t-1} | y_{1:t-1}) d\mathbf{x}_t dv_{t-1} d\lambda_{t-1} \end{aligned} \tag{5}$$

because

$$\iint_{\mathbb{R}_+ \times \mathbb{R}} u(v_{t-1}, \lambda_{t-1}, j_{t-1}^y, j_{t-1}^v | y_{1:t-1}) dj_{t-1}^y dj_{t-1}^v = u(v_{t-1}, \lambda_{t-1} | y_{1:t-1}).$$

---

<sup>5</sup>Being nested models, the SV, SVYJ and SVCJ models could also be estimated with the same method by reducing the problem's dimension.

Conditioning on the number of jumps also allows us to further simplify Equation (5) so that

$$f(y_t | y_{1:t-1}) = \sum_{n=0}^{\infty} \int \cdots \int_{\mathbb{R}_+^5} \mathbb{P}(n_t | \lambda_{t-1}) r(y_t | v_t, v_{t-1}, \lambda_t, \lambda_{t-1}, j_t^v, n_t = n) \\ q(v_t | v_{t-1}, j_t^v) q(j_t^v | n_t = n) q(\lambda_t | \lambda_{t-1}) \\ u(v_{t-1}, \lambda_{t-1} | y_{1:t-1}) dv_t d\lambda_t dj_t^v dv_{t-1} d\lambda_{t-1}, \quad (6)$$

as

$$\int_{\mathbb{R}} r(y_t | v_t, v_{t-1}, \lambda_t, \lambda_{t-1}, j_t^y, j_t^v, n_t = n) q(j_t^y | n_t = n) dj_t^y = r(y_t | v_t, v_{t-1}, \lambda_t, \lambda_{t-1}, j_t^v, n_t = n),$$

which has a normal density. Specifically, we have that

$$y_t | v_t, v_{t-1}, \lambda_t, \lambda_{t-1}, j_t^v, n_t = n \sim \mathcal{N}(\mu_t, \sigma_t^2),$$

where

$$\mu_t = \left( \mu - \frac{1}{2}v_{t-1} - \bar{\alpha}\lambda_{t-1} \right) h + \sqrt{v_{t-1}h} (\rho_v \varepsilon_t^v + \rho_\lambda \varepsilon_t^\lambda) + \alpha n + \rho_z j_t^v, \\ \sigma_t^2 = v_{t-1} (1 - \rho_v^2 - \rho_\lambda^2) h + n \delta^2, \\ \varepsilon_t^v = \frac{v_t - v_{t-1} - \kappa(\theta - v_{t-1})h - j_t^v}{\sigma \sqrt{v_{t-1}h}} \quad \text{and} \quad \varepsilon_t^\lambda = \frac{\lambda_t - \lambda_{t-1} - \chi(\theta - \lambda_{t-1})h}{\xi \sqrt{\lambda_{t-1}h}}.$$

We can obtain the integrand of Equation (6) in closed form because

$$n_t | \lambda_{t-1} \sim \mathcal{P}(\lambda_{t-1}h), \\ v_t | v_{t-1}, j_t^v \sim \mathcal{N}(v_{t-1} + \kappa(\theta - v_{t-1})h + j_t^v, \sigma^2 v_{t-1}h), \\ j_t^v | n_t = n \sim \Gamma(n, \nu), \\ \lambda_t | \lambda_{t-1} \sim \mathcal{N}(\lambda_{t-1} + \chi(\omega - \lambda_{t-1})h, \xi^2 \lambda_{t-1}h)$$

and the corresponding densities are known in closed form as well. Note that  $\mathcal{P}(m)$  is a Poisson distribution with mean  $m$  and  $\Gamma(n, \nu)$  is a gamma distribution with mean  $n\nu$ .

Finally, one can get the time- $t$  posterior *density* of the latent states in a similar fashion by computing the following integral

$$u(\mathbf{x}_t | y_{1:t}) = u(v_t, \lambda_t | y_{1:t}) \\ = \frac{1}{f(y_t | y_{1:t-1})} \left( \sum_{n=0}^{\infty} \iiint_{\mathbb{R}_+^3} \mathbb{P}(n_t | \lambda_{t-1}) r(y_t | v_t, v_{t-1}, \lambda_t, \lambda_{t-1}, j_t^v, n_t = n) \\ q(v_t | v_{t-1}, j_t^v) q(j_t^v | n_t = n) q(\lambda_t | \lambda_{t-1}) u(v_{t-1}, \lambda_{t-1} | y_{1:t-1}) dj_t^v dv_{t-1} d\lambda_{t-1} \right).$$

### 3.3 Numerical implementation

To be able to numerically compute the integrals above, we need to define a discretization of the state space—so-called *likely* sequences. The interval boundaries used to discretize the state space of the variance factor  $v$ , the jump intensity factor  $\lambda$  and the variance jumps  $j$ , are chosen over the following ranges:

$$V = [v^{(1)} \quad \cdots \quad v^{(N)}], \quad \Lambda = [\lambda^{(1)} \quad \cdots \quad \lambda^{(M)}], \quad J = [j^{(1)} \quad \cdots \quad j^{(K)}], \\ v^{(1)} = \mathcal{E}_\infty^V - \delta_N^V \sqrt{\mathcal{V}_\infty^V}, \quad \lambda^{(1)} = \mathcal{E}_\infty^\Lambda - \delta_M^\Lambda \sqrt{\mathcal{V}_\infty^\Lambda}, \quad j^{(1)} = \mathcal{E}_\infty^J - \delta_K^J \sqrt{\mathcal{V}_\infty^J}, \\ v^{(N)} = \mathcal{E}_\infty^V + \delta_N^V \sqrt{\mathcal{V}_\infty^V}, \quad \lambda^{(M)} = \mathcal{E}_\infty^\Lambda + \delta_M^\Lambda \sqrt{\mathcal{V}_\infty^\Lambda}, \quad j^{(K)} = \mathcal{E}_\infty^J + \delta_K^J \sqrt{\mathcal{V}_\infty^J},$$

$$\begin{aligned} \mathcal{E}_\infty^V &= \lim_{t \rightarrow \infty} \mathbb{E}^\mathbb{P} [v_t | \mathcal{F}_0], & \mathcal{E}_\infty^\Lambda &= \lim_{t \rightarrow \infty} \mathbb{E}^\mathbb{P} [\lambda_t | \mathcal{F}_0], & \mathcal{E}_\infty^J &= \lim_{t \rightarrow \infty} \mathbb{E}^\mathbb{P} [j_t^v | \mathcal{F}_0], \\ \mathcal{V}_\infty^V &= \lim_{t \rightarrow \infty} \text{Var}^\mathbb{P} [v_t | \mathcal{F}_0], & \mathcal{V}_\infty^\Lambda &= \lim_{t \rightarrow \infty} \text{Var}^\mathbb{P} [\lambda_t | \mathcal{F}_0], & \mathcal{V}_\infty^J &= \lim_{t \rightarrow \infty} \text{Var}^\mathbb{P} [j_t^v | \mathcal{F}_0], \end{aligned}$$

where  $\mathcal{E}_\infty^V$  and  $\mathcal{V}_\infty^V$  are the long-run expected value and variance of the stationary process of the variance factor, respectively. Moreover,  $\mathcal{E}_\infty^\Lambda$ ,  $\mathcal{V}_\infty^\Lambda$ ,  $\mathcal{E}_\infty^J$  and  $\mathcal{V}_\infty^J$  are defined analogously for the jump intensity factor and for the variance jumps, respectively. The intermediate points are determined by the following equations:

$$\begin{aligned} v^{(i)} &= \left( \sqrt{v^{(1)}} + \left( \frac{i-1}{N-1} \right) \left( \sqrt{v^{(N)}} - \sqrt{v^{(1)}} \right) \right)^2, & i &= 2, \dots, N-1, \\ \lambda^{(l)} &= \left( \sqrt{\lambda^{(1)}} + \left( \frac{l-1}{M-1} \right) \left( \sqrt{\lambda^{(M)}} - \sqrt{\lambda^{(1)}} \right) \right)^2, & l &= 2, \dots, M-1, \\ j^{(j)} &= j^{(1)} + \left( \frac{j-1}{M-1} \right) \left( j^{(K)} - j^{(1)} \right), & j &= 2, \dots, K-1, \end{aligned}$$

making the variance (jump intensity) likely sequences uniformly distributed in the volatility (square root of the jump intensity) domain.

The functions  $\delta_N^V$ ,  $\delta_M^\Lambda$  and  $\delta_K^J$  assure that the DNF converges and shall therefore respect two sets of conditions:

$$\begin{aligned} \lim_{N \rightarrow \infty} \delta_N^V &= \infty, & \lim_{M \rightarrow \infty} \delta_M^\Lambda &= \infty, & \lim_{K \rightarrow \infty} \delta_K^J &= \infty, \\ \lim_{N \rightarrow \infty} \frac{\delta_N^V}{N} &= 0, & \lim_{M \rightarrow \infty} \frac{\delta_M^\Lambda}{M} &= 0, & \lim_{K \rightarrow \infty} \frac{\delta_K^J}{K} &= 0. \end{aligned}$$

The first set of conditions makes sure that we cover the domain as the number of nodes increases. On the other hand, the second set of conditions ensures that the partition becomes finer and finer as the number of nodes increases. Obviously, a myriad of such functions exists. In this study, we use  $\delta_N^V = 3 + \log(N)$ ,  $\delta_M^\Lambda = 3 + \log(M)$  and  $\delta_K^J = 3 + \log(K)$ , although other functions could be used (as long as they satisfy the two sets of conditions).

Moreover, in the spirit of Langrock et al. (2012), a special integration rule will be used in this study: broadly speaking, if  $g_1$  and  $g_2$  are two functions, then

$$\int_a^b g_1(x)g_2(x) dx \approx g_1(c) \int_a^b g_2(x) dx, \tag{7}$$

where  $c$  is a representative point in  $[a, b]$  (e.g., the midpoint). Given that we are integrating products of probability density functions (pdfs), this integral rule will allow us to rewrite the integral of Equation (7) as  $g_1(c)(G_2(b) - G_2(a))$ , where  $G_2$  is the corresponding cumulative distribution function of the pdf  $g_2$ .<sup>6</sup>

Then, to compute numerical integrals, it is often more convenient to define intervals. Therefore, let

$$\begin{aligned} V^{(i)} &= \left[ \frac{v^{(i-1)} + v^{(i)}}{2}, \frac{v^{(i)} + v^{(i+1)}}{2} \right), & i &= 1, \dots, N, \\ \Lambda^{(l)} &= \left[ \frac{\lambda^{(l-1)} + \lambda^{(l)}}{2}, \frac{\lambda^{(l)} + \lambda^{(l+1)}}{2} \right), & l &= 1, \dots, M, \\ J^{(j)} &= \left[ \frac{j^{(j-1)} + j^{(j)}}{2}, \frac{j^{(j)} + j^{(j+1)}}{2} \right), & j &= 1, \dots, K, \end{aligned}$$

<sup>6</sup>Other quadratures have been proposed in the literature. For instance, Fridman and Harris (1998) use of Gauss-Legendre quadrature. Bartolucci and De Luca (2001) and Clements et al. (2006a) use a simple numerical integral scheme based on the midpoint  $c$ , i.e.,

$$\int_a^b f_1(x)g_2(x) dx \approx (b-a)g_1(c)g_2(c).$$

be three different sets of intervals that cover  $[0, \infty)$ , where  $v^{(0)} = -v^{(1)}$ ,  $\lambda^{(0)} = -\lambda^{(1)}$ ,  $j^{(0)} = -j^{(1)}$ ,  $v^{(N+1)} = \infty$ ,  $\lambda^{(M+1)} = \infty$ , and  $j^{(K+1)} = \infty$ .

Finally, the time- $t$  likelihood contribution is approximated by

$$\begin{aligned} \hat{f}(y_t | y_{1:t-1}) \approx & \sum_{n=0}^R \sum_{i_t=1}^N \sum_{i_{t-1}=1}^N \sum_{j_t=1}^K \sum_{l_t=1}^M \sum_{l_{t-1}=1}^M r(y_t | v^{(i_t)}, v^{(i_{t-1})}, \lambda^{(l_t)}, \lambda^{(l_{t-1})}, j^{(j_t)}, n_t = n) \\ & \mathbb{P}(n_t = n | \lambda^{(l_{t-1})}) q(v_t \in V^{(i_t)} | v^{(i_{t-1})}, j^{(j_t)}) \\ & q(j_t^v \in J^{(j_t)} | n_t = n) q(\lambda_t \in \Lambda^{(l_t)} | \lambda^{(l_{t-1})}) \\ & u(v^{(i_{t-1})}, \lambda^{(l_{t-1})} | y_{1:t-1}), \end{aligned} \quad (8)$$

where  $R$  is the truncation level of the Poisson random variable. The time- $t$  posterior *density*  $u(v_t, \lambda_t | y_{1:t})$  is also estimated in a similar way:

$$\begin{aligned} \hat{u}(v^{(i_t)}, \lambda^{(l_t)} | y_{1:t}) \approx & \frac{1}{\hat{f}(y_t | y_{1:t-1})} \sum_{n=0}^R \sum_{i_t=1}^N \sum_{j_t=1}^K \sum_{l_t=1}^M r(y_t | v^{(i_t)}, v^{(i_{t-1})}, \lambda^{(l_t)}, \lambda^{(l_{t-1})}, j^{(j_t)}, n_t = n) \\ & \mathbb{P}(n_t = n | \lambda^{(l_{t-1})}) q(v_t \in V^{(i_t)} | v^{(i_{t-1})}, j^{(j_t)}) \\ & q(j_t^v \in J^{(j_t)} | n_t = n) q(\lambda_t \in \Lambda^{(l_t)} | \lambda^{(l_{t-1})}) \\ & u(v^{(i_{t-1})}, \lambda^{(l_{t-1})} | y_{1:t-1}). \end{aligned}$$

## 4 Accuracy and reliability

We now aim to assess the accuracy and reliability of the DNF. Specifically, we first verify that it converges to the correct likelihood value by comparing the log-likelihood obtained with the DNF to that of the SIR assuming that a large computing budget is available. This experiment will tell us how the DNF performs in a best-case scenario.

As explained in Malik and Pitt (2011), the SIR-based likelihood function is unbiased (see Section A of the Supplementary Material for more details on the SIR method). Therefore, when a large computing budget is used, the SIR-based likelihood estimate should be very close to its true value with very little sampling error. Hence, it can be used as a reliable benchmark to assess the accuracy and reliability of the DNF.

For the SIR method, we use an increasing number of particles as a function of the model complexity. Specifically, we use 100,000 particles for the SV model, 250,000 particles for the SVYJ model and 1,000,000 particles for both the SVCJ and SVCJSI models. Then, we set the number of nodes in the variance (variance jump) grid to  $N = 200$  ( $K = 100$ ) for the SV, SVYJ and SVCJ models.<sup>7</sup> As for the SVCJSI model, we set  $N = M = 50$  and  $K = 25$  due to memory size limits.<sup>8</sup>

### 4.1 Random series

In this first test, we compare the likelihood of 1,000 one-year daily series taken at random using parameters that are also randomly generated. The aim of this test is to determine the accuracy of the DNF for a wide range of possible paths and parameter sets.

<sup>7</sup>According to our results, a reliable rule of thumb is to set the number of nodes of the variance jump grid,  $K$ , to be two and three times fewer than the number of points of the variance.

<sup>8</sup>We used a graphics processing unit (GPU) to compute the likelihood function of the SVCJSI model with the DNF; see Section 5.1.

Therefore, we generate parameter sets using the following bounds:

$$\begin{aligned}
-0.20 &\leq \mu \leq 0.20 \\
0.00 &\leq \kappa \leq 10.00 \\
0.00 &\leq \theta \leq 0.10 \\
0.10 &\leq \sigma \leq 1.00 \\
-0.95 &\leq \rho_v \leq 0.95 \\
0.00 &\leq \chi \leq 50.00 \\
0.00 &\leq \omega \leq 25.00 \\
0.10 &\leq \xi \leq 10.00 \\
-\max[0.95, 1 - \rho_v^2] &\leq \rho_\lambda \leq \max[0.95, 1 - \rho_v^2] \\
-0.05 &\leq \alpha \leq 0.05 \\
0.00 &\leq \delta \leq 0.10 \\
-5.00 &\leq \rho_z \leq 5.00 \\
0.00 &\leq \nu \leq 0.03.
\end{aligned}$$

These values are reasonable bounds for the parameters; they span multiple potential and realistic cases. Then, based on the simulated parameters, we generate one-year paths and we compute the log-likelihood using both methods. Specifically, we compare both estimates using the absolute percentage error (APE):

$$\text{APE} = 100\% \left| \frac{\log(\mathcal{L}_{\text{DNF}}(\Theta)) - \log(\mathcal{L}_{\text{SIR}}(\Theta))}{\log(\mathcal{L}_{\text{SIR}}(\Theta))} \right|. \quad (9)$$

Table 1 reports seven quantiles of the APE distribution for the four models considered. Broadly speaking, the medians hover between 0.01% and 0.05% for the four models which is very small. The right tail of the APE distributions is also very thin with 99.5<sup>th</sup> quantiles that are between 0.40% and 0.85%.

**Table 1: Distribution of the log-likelihood absolute percentage error: Random series.**

$\alpha$	SV	SVYJ	SVCJ	SVCJSI
0.250	0.0039	0.0048	0.0133	0.0159
0.500	0.0088	0.0123	0.0512	0.0367
0.750	0.0168	0.0263	0.1619	0.0843
0.900	0.0318	0.0530	0.3608	0.1931
0.950	0.0545	0.0781	0.5016	0.3773
0.990	0.1986	0.2108	0.6838	0.7536
0.995	0.4318	0.5068	0.7235	0.8247

This table reports the quantiles of the absolute percentage error distribution for the four models considered in this study. These distributions are obtained by generating 1,000 one-year series using randomly selected parameters and by computing the APE for each path as shown in Equation (9). SV stands for stochastic volatility, SVYJ for stochastic volatility with return jumps, SVCJ for stochastic volatility with correlated and simultaneous jumps in return and variance, and SVCJSI for SVCJ with stochastic jump intensity.

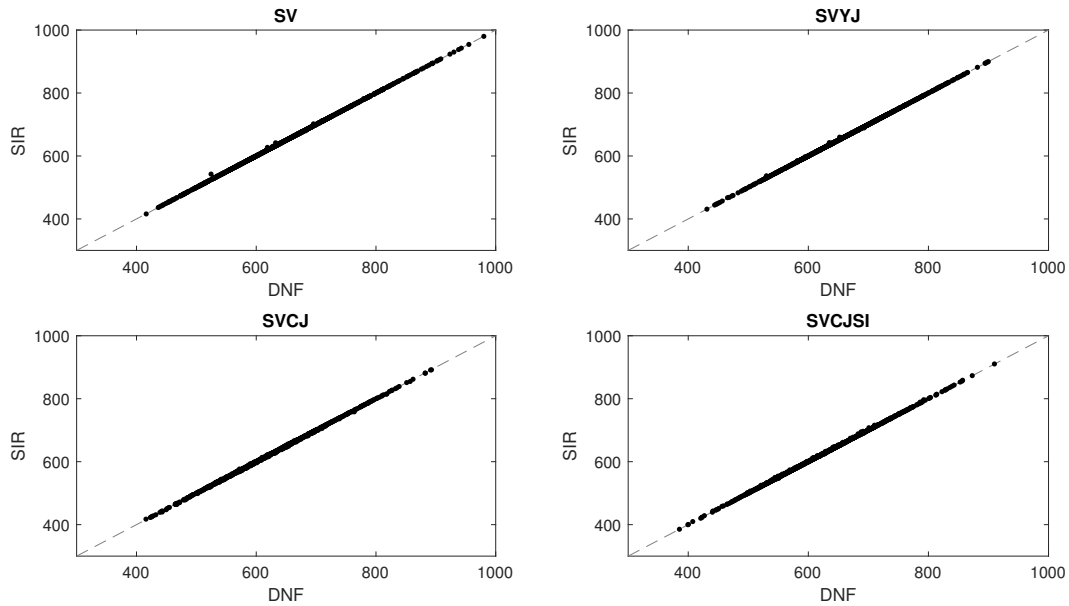
Figure 1 supplements Table 1 by reporting scatter plots of SIR- and DNF-based log-likelihoods. Most of these are aligned on the diagonal meaning that both methods are yielding very much alike estimates when using a large computing budget. Based upon this first test, we can comfortably conclude that the DNF is very accurate, at least within the range of parameters provided above and for a large computing budget.

## 4.2 S&P 500 Index

Instead of using randomly generated series that might not resemble those used in empirical finance, we now focus on the S&P 500 returns as these would be representative of typical asset returns. Specifically, we use daily S&P 500 Index returns (excluding dividends), from January 1990 to September 2018,

obtained from the Bloomberg terminal. For this specific time series, we then randomly generate a set of parameters and compute the resulting likelihood with both the SIR and DNF methods.

Table 2 mimics the results of Table 1, only this time we consider the S&P 500 returns over the last three decades. The results are similar to those obtained with random series: in fact, the median APE is rather small, ranging between 0.02% and 0.10%. Again, even the worst cases are still very decent, with the 99.5<sup>th</sup> quantile being below 1% for the four models. Therefore, for a typical financial time series and a large computing budget, the DNF is very accurate.



**Figure 1: SIR- against DNF-based log-likelihood estimates: Random series.** This figure reports the SIR- and DNF-based log-likelihood estimates for large computing budget and for the four models considered in this study. These scatter plots are obtained by generating 1,000 one-year series using randomly selected parameters and by computing the log-likelihood for each path by using both the DNF and the SIR methods.

**Table 2: Distribution of the log-likelihood absolute percentage error: S&P 500 Index.**

$\alpha$	SV	SVYJ	SVCJ	SVCJSI
0.250	0.0118	0.0161	0.0122	0.0403
0.500	0.0245	0.0358	0.0311	0.1020
0.750	0.0513	0.0712	0.0659	0.2305
0.900	0.1049	0.1155	0.1152	0.4243
0.950	0.1793	0.2025	0.1477	0.5971
0.990	0.3123	0.5546	0.2991	0.7696
0.995	0.3503	0.6336	0.4236	0.7812

This table reports the quantiles of the absolute percentage error distribution for the four models considered in this study. These distributions are obtained by generating 1,000 sets of parameters and by computing the APE for each set as shown in Equation (9) with the S&P 500 Index time series.

## 5 Precision and computing times

As shown in the previous section, the DNF is capable of providing very accurate log-likelihood estimates when a large computing budget is available. Nonetheless, it is of great interest to also assess its performance when a smaller budget is considered because most inference methods—both frequentist and Bayesian—require a considerable number of likelihood function evaluations. The overarching goal of this section is to analyze the precision of the DNF compared with computing times required to attain such precision.

We first discuss the use of central processing unit (CPU) and graphics processing unit (GPU) for the problem at hand. We then compare the DNF and SIR in terms of computing times and precision. Finally, we analyze the speed and accuracy trade-off of the DNF by determining the grid size or computing time necessary to achieve a given level of precision (and vice versa).

## 5.1 Motivation for CPU and GPU computing


The computing times involved in this study are greatly influenced by the number of state variables considered, which has an impact on the dimension of the sum of Equation (8). Indeed, Equation (8) involves a sum over two, three, four or six dimensions for the SV, SVYJ, SVCJ and SVCJSI models, respectively. Obviously, such calculations can become cumbersome when  $N$ ,  $M$ ,  $K$ ,  $R$  and  $T$  are large.

It is easy to see why a GPU can be very useful. For instance, if  $N = 50$ ,  $M = 50$ ,  $K = 25$  and  $R = 2$ , then the size of the six-dimensional matrix containing all the elements that need to be summed up is

$$50 \times 50 \times 50 \times 50 \times 25 \times 2 \times 8 \text{ bytes} = 2,500,000,000 \text{ bytes} = 2.5 \text{ gigabytes.}$$

Moreover, the number of operations needed to obtain one likelihood contribution is rather large and it increases exponentially, e.g.,  $5 \times 50 \times 50 \times 50 \times 50 \times 25 \times 2 = 1,562,500,000$  multiplications, and  $50 \times 50 \times 50 \times 50 \times 25 \times 2 = 312,500,000$  additions for the example above.

CUDA-enabled GPUs are extremely efficient when it comes to calculating sums. In fact, for very large matrices, the time needed to compute sums is manifold lower when using a GPU instead of a CPU. Figure 2 shows the computing time for summing all the entries of an  $n \times n$  matrix while using a GPU (black) or a CPU (grey). For instance, the GPU-based calculation is eight times faster for a matrix of size  $20,000 \times 20,000$ —which is a significant decrease.



**Figure 2: Computing time against matrix size: GPU versus CPU.** This figure reports the time needed to compute the sum of all the entries of an  $n \times n$  matrix, where  $n$  is the matrix size, as a function of a matrix size. Central processing unit is denoted by CPU and graphics processing unit by GPU.

But for GPU computing to cut overall computing times, it is important to take into account the additional time required to move data between the CPU and GPU, known as the overhead. For array sizes involved with deterministic jump arrival intensity models, our experiments showed that a conventional CPU is still faster overall.

Therefore and unless stated otherwise, the SIR and DNF methods have been coded in Matlab 2018b on a single thread with a typical CPU. However, the DNF method applied to the SVCJSI model has been implemented with an NVIDIA QUADRO P6000 GPU, also using Matlab 2018b.<sup>9</sup>

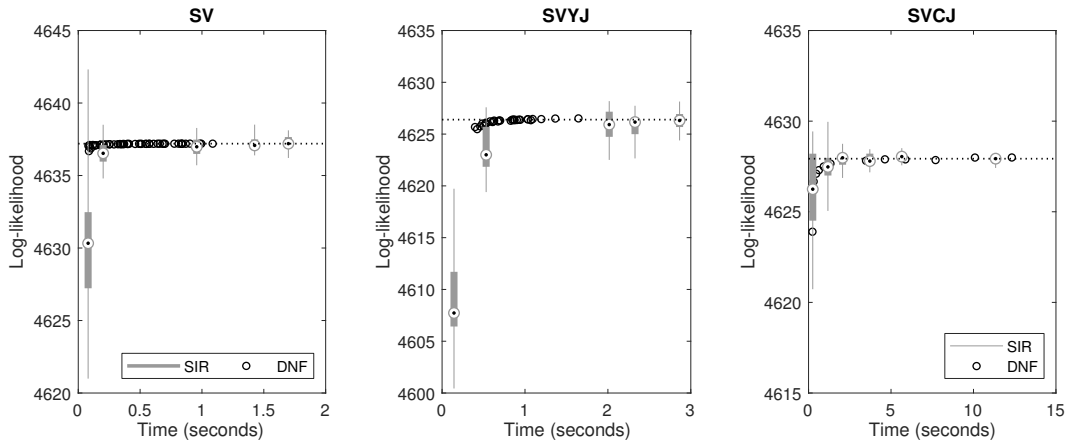
## 5.2 Deterministic jump arrival intensity: SV, SVYJ and SVCJ models

Since models with deterministic and stochastic jump arrival intensity have been implemented differently, we analyze computing times and precision for both families of models separately.

### 5.2.1 Computing times: Comparison between DNF and SIR

The goal of this first test is to analyze the accuracy of SIR and DNF for different fixed and finite computing budgets. For the SIR method, the computing budget is determined by the number of particles whereas for the DNF method, the grid size parameters  $N$  and  $K$  determine the computing budget. For all computations, we use the same daily S&P 500 return series over 5 years (about 1250 observations) evaluated with the parameters that produced the highest likelihood in Section 4.2 as a proxy for maximum likelihood parameters.

Then, for a given computing budget, we compute the likelihood value and record the computing time. Moreover, for the SIR, we repeat the exercise 50 times to approximate the distribution of the log-likelihood values. This entire exercise is then rerun for several computing budgets.



**Figure 3: DNF-based log-likelihood estimates and the distribution of the SIR-based log-likelihood estimates against time.** This figure reports the DNF-based log-likelihood estimates (black circles) and the distribution of the SIR-based log-likelihood estimates (grey box plots), while using small computing budgets. We use the same daily S&P 500 return series over a five year period. For the SIR, we repeat the exercise 50 times to get a distribution of the log-likelihood value. The size of the budget is proxied by the computing time—the number of seconds it took to run one log-likelihood evaluation. We consider boxplots with whiskers from minimum to maximum and boxes from the lower quartile to the upper one. A circle with a dot represents the median of the SIR distribution.

Figure 3 reports this comparison. We show boxplots (in grey) of the distribution of the SIR-based log-likelihood values and we compare them to values obtained with the DNF method (black circles). The *true* value (dashed line) is given by the average obtained with SIR using a very large computing budget (this number is virtually the same as that obtained with the DNF method for a large budget as well).

For the SV model, the DNF method yields values that are very close to the true log-likelihood value—and this, after only 100 milliseconds (ms). For such small computing budget, the SIR does not produce adequate results: it is biased and shows important variations. For a budget of one second per

<sup>9</sup>The computing times are obtained by running Matlab programs on a desktop computer with two 2.1 GHz Intel Xeon E5-2620 v4 and 48 GB RAM. The NVIDIA QUADRO P6000 we use has 3840 CUDA parallel-processing cores and a GPU memory of 24 GB.

likelihood valuation, the SIR still gives very noisy log-likelihood estimates even though this value is asymptotically consistent: the whiskers are quite long, but the median is aligned with the true value.

The DNF provides similar results for the SVYJ model: after 600 ms, the deterministic method produces very steady and stable estimates of the log-likelihood function—consistent with the true value obtained with a large computing budget. The SIR-based log-likelihood, on the other hand, is still volatile even with a budget of two seconds.

The results for the SVCJ model tell a similar story: after a few seconds, the DNF is precise and stable and quickly converges to the true value whereas the SIR-based log-likelihood still shows significant sampling variations as the whiskers are still clearly visible to the naked eye.

### 5.2.2 Speed and accuracy trade-off

We know the DNF method performs very well compared to the SIR providing accurate results much quicker than the SIR. In this second test, we would like to determine the precision gained (or lost) if we have a limited computing budget, as defined by either the grid size or computing time.

In this experiment, we generate 100 sets of parameters and we evaluate the likelihood function for varying number of nodes. We use the number of variance nodes  $N$  ranging from 25 to 200 with steps of 5. The number of variance jump nodes is consistent with the rule of thumb explained above, i.e.,  $N/2.5$ . The number of nodes associated with the jumps,  $R$ , is set to 2. We still use daily S&P 500 returns over 5 years for about 1250 observations.

The likelihood values obtained for these varying number of nodes is then compared to that obtained with the SIR while using a large computing budget (see Section 4). We still use the APE as a basis of comparison and compute mean APE (MAPE) to assess the average difference between the log-likelihood values:

$$\text{MAPE} = 100\% \left( \frac{1}{100} \sum_{i=1}^{100} \left| \frac{\log(\mathcal{L}_{\text{DNF}}(\Theta_i)) - \log(\mathcal{L}_{\text{SIR}}(\Theta_i))}{\log(\mathcal{L}_{\text{SIR}}(\Theta_i))} \right| \right).$$

The leftmost panels of Figure 4 report the smoothed version of the MAPE against size—here, proxied by the number of nodes in the variance grid  $N$ .<sup>10</sup> For the four models considered, a size between 50 and 60 would lead to a MAPE below 0.1%. This result is in fact consistent with the size used in Watanabe (1999), Clements et al. (2004), Clements et al. (2006a) and Langrock et al. (2012).

The rightmost panels of Figure 4 detail the relationship between the MAPE and the computing time per likelihood valuation (in seconds). Again, as explained above, the computing time tends to increase with the complexity of the model. It takes approximately 75 ms, 200 ms and 300 ms to attain an MAPE of about 0.1% with the SV, SVYJ and SVCJ models, respectively. Therefore, with a conventional CPU, one is able to compute the likelihood function in much less than a second and reach an MAPE of 0.1%.

### 5.3 Stochastic jump arrival intensity: SVCJSI model

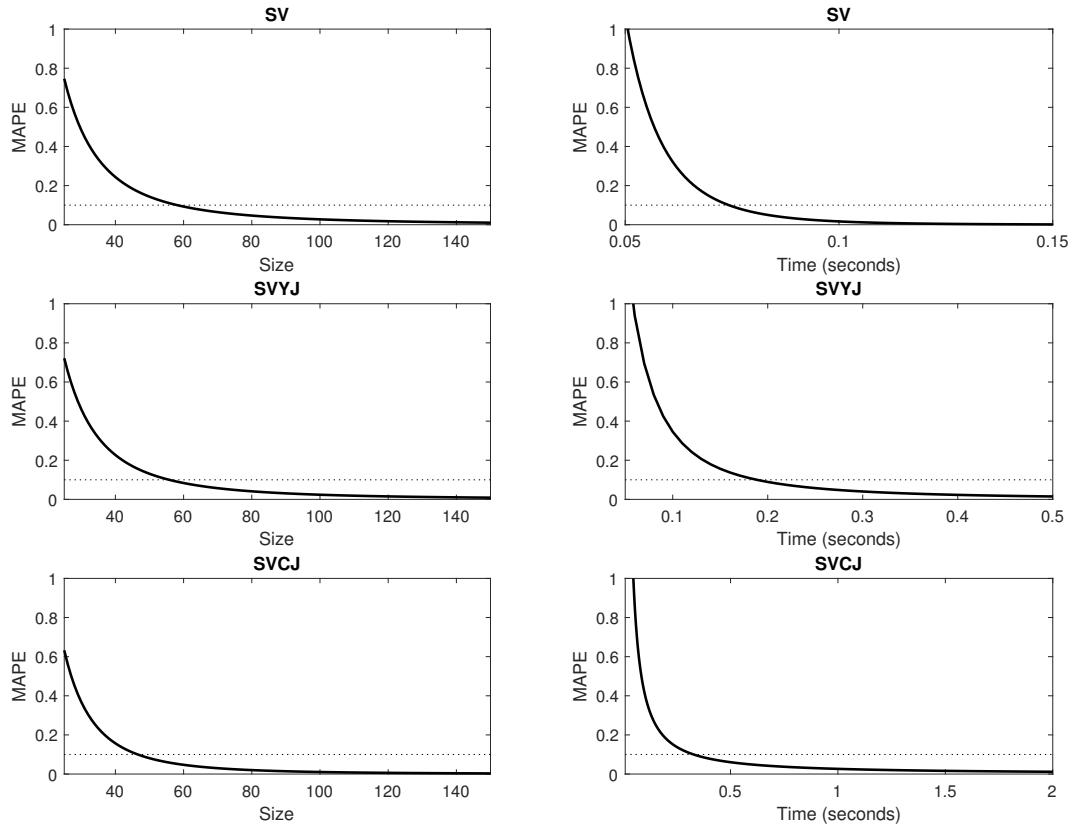
We now analyze the performance of the DNF for the SVCJSI model. As opposed to the family of deterministic jump arrival intensity models, the SVCJSI model is comprised of four state variables and to accelerate computations, we implemented the DNF using a high-end GPU.

As in Section 5.2, we first seek to compare the precision and computing times of the SIR and DNF methods. Figure 5 is the same as Figure 3, but for the SVCJSI model.

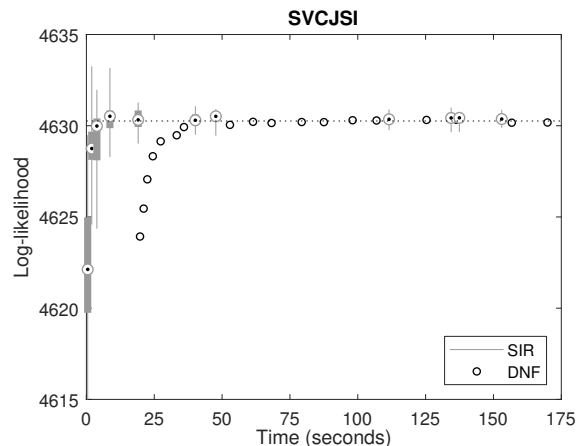
For a very small computing budget (a couple of seconds), the SIR generates a very wide range of likelihood values. As we add particles and thus the required computing time, the width of whiskers

<sup>10</sup>We use Matlab's `fit` function with a power model, i.e.,  $y = ax^b$ , to fit the MAPE curve. The results are robust to other specifications.

decreases as expected. However, even if the computing budget is large (2 or 3 minutes per likelihood valuation), there is still a significant uncertainty that remains around the likelihood value obtained. But if we are equipped with a powerful GPU, one is capable of achieving a very precise likelihood value with the DNF within a minute.



**Figure 4: Mean absolute percentage error against the number of variance nodes and the computing time.** This figure reports the smoothed MAPE per model against the size—proxied as the number of variance nodes—on the leftmost panels and against computing time (in seconds) on the rightmost panels. We use daily returns from the S&P 500 over 5 years, generate 100 sets of parameters and we evaluate the likelihood function for varying number of nodes. We use number of variance nodes  $N$  ranging from 25 to 200 with steps of 5. The number of variance jump nodes is consistent with the rule of thumb explained above, i.e.,  $N/2.5$ . The number of nodes associated with the jumps,  $R$ , is set to 2 in all of our tests.




**Figure 5: DNF-based log-likelihood estimates and the distribution of the SIR-based log-likelihood estimates against time, continued.** See the description of Figure 3 for more details.

Figure 5 also illustrates the difficulty of precisely evaluating the likelihood function for the SVCJSI model (with the technology available to most users in 2019). If the computing budget available is small (e.g., a couple of seconds), then one should turn to the SIR implemented on a CPU but the resulting likelihood value will be biased and random. If a GPU is available, the computing budget must be large enough to overcome the extra overhead: at least 30 to 60 seconds are needed to achieve a level of accuracy that is impossible to attain with the SIR with a finite computing budget.

We also analyze the MAPE as a function of the grid size or computing time. The number of variance and jump intensity nodes ranges from 25 to 50 with steps of 1, with  $N = M$ .<sup>11</sup> Moreover, the number of variance jumps is set to  $\lceil N/2.5 \rceil$  and the number of jump nodes is set to 2 still.

The results are illustrated in Figure 6 for the SVCJSI model. Again, we observe that a grid size of about 50-60 achieves an MAPE of 0.1% (as with the SV, SVYJ and SVCJ models) but the required computing time to achieve such MAPE is about 160 seconds with GPU computing.



**Figure 6: Mean absolute percentage error against the number of variance nodes and the computing time, continued.** This figure reports the smoothed MAPE per model against the size—proxied as the number of variance nodes—on the leftmost panels and against computing time (in seconds) on the rightmost panels. We use daily returns from the S&P 500 over 5 years and generate 100 sets of parameters and we evaluate the likelihood function for varying number of nodes. For the SVCJSI model, the number of variance and jump intensity nodes ranges from 25 to 50 with steps of 1, with  $N = M$ . The number of variance jumps is set to  $\lceil N/2.5 \rceil$  and the number of jump nodes is set to 2 still.


## 6 Empirical applications

In the previous applications, we have shown the precision of the DNF in evaluating the likelihood function. For frequentist-based inference such as the maximum likelihood, the likelihood function must be smooth enough for an optimizer to find a global maximum in the parameter space. As such, the DNF has a key advantage over the SIR as its construction leads to a smooth likelihood function.

We thus present in this last section an important application of the DNF methodology: we compute the maximum likelihood estimates (MLE) of each of the four models studied in this paper. This section is certainly one of the few attempts in the literature to compute the MLE of such complex financial models.

We use the S&P 500 return series (excluding dividends) from 1990 to 2018 (see Figure 7). This time frame contains a few recessions and eras of turmoil: the early 1990s recession (subsequent to the 1990 oil price shock), the early 2000s recession (partly caused by the collapse of the speculative dot-com bubble and the 9/11 attacks) and the Great Recession (a by-product of the collapse of the US housing bubble, followed by a global financial crisis). This sample should therefore contain a time-varying volatility, return jumps, variance jumps, and potentially, a changing jump intensity. The grid sizes used for the computing of the likelihood function are such that  $N = M = 50$ ,  $K = 20$  and  $R = 2$ , which are consistent with our results.

<sup>11</sup>A number of nodes of 50 is the physical size limit the GPU handles at the time of writing this article.



**Figure 7: S&P 500 Index returns (1990–2018).** This figure reports daily S&P 500 Index returns (excluding dividends), from January 1990 to September 2018, obtained from the Bloomberg terminal.

Table 3 reports the MLE for each model along with their standard error in brackets.<sup>12</sup> These estimates are comparable to those found in other studies and using different statistical methodologies (e.g., Andersen et al., 2002; Eraker et al., 2003; Christoffersen et al., 2010; Hurn et al., 2015).

**Table 3: Estimated parameters for the S&P 500 Index.**

	SV	SVYJ	SVCJ	SVCJSI
$\mu$	0.041 (0.017)	0.035 (0.016)	0.038 (0.020)	0.035 (0.024)
$\kappa$	5.923 (0.405)	6.357 (0.343)	3.689 (0.311)	4.316 (0.201)
$\theta$	0.031 (0.002)	0.027 (0.001)	0.032 (0.002)	0.034 (0.002)
$\sigma$	0.514 (0.017)	0.488 (0.010)	0.446 (0.018)	0.452 (0.001)
$\rho_v$	-0.692 (0.024)	-0.708 (0.024)	-0.745 (0.024)	-0.666 (0.031)
$\chi$	–	–	–	2.706 (2.618)
$\omega$	–	2.487 (1.487)	5.125 (3.021)	3.232 (3.111)
$\xi$	–	–	–	6.947 (5.778)
$\rho_\lambda$	–	–	–	-0.411 (0.102)
$\alpha$	–	-0.014 (0.007)	-0.007 (0.005)	-0.014 (0.016)
$\delta$	–	0.008 (0.003)	0.003 (0.005)	0.005 (0.013)
$\nu$	–	–	0.004 (0.001)	0.011 (0.006)
$\rho_z$	–	–	-1.809 (0.671)	-1.381 (0.706)

This table reports the maximum likelihood estimates for the four models considered in this study. We use the S&P 500 return series from 1990 to 2018. Robust standard errors (in brackets) are calculated from the outer product of the gradient at the optimum parameter value.

Figure 8 shows the filtered annualized volatility (i.e.,  $\sqrt{V_t/252}$ ), the jump intensity, the return jumps and the variance jumps. Generally speaking, the volatility is rather similar between the various models, the only difference being that the SVCJ and SVCJSI models allow for sharper increases because of the volatility jumps. For instance, during the last financial crisis, the volatility goes higher with the models that permit variance jumps.

<sup>12</sup>Robust standard errors (in brackets) are calculated from the outer product of the gradient at the optimum parameter value. For more details on this method, see Appendix B.5 of Rémillard (2013).

For models with constant jump intensity, the number of jumps per year is small—about three to five, on average. This figure is consistent with other studies such as Eraker et al. (2003) and Hurn et al. (2015). The filtered return jumps are rather similar for the SVYJ and SVCJ models.


On the other hand, the SVCJSI model allows for jump clustering as the jump intensity varies over time, and the jump processes are slightly different of those obtained with the SVYJ and SVCJ models. Specifically, the number of filtered jumps reduces during calm eras and increases during crises.

## 7 Concluding remarks

In this paper, we developed a new estimation method for jump-diffusion models based upon the DNF method of Kitagawa (1987). The method we introduced allows for the likelihood evaluation of models that include stochastic volatility, return jumps, variance jumps, as well as stochastic jump intensity.

We then compared the performance of the DNF with the SIR. We found that the proposed DNF method is stable, accurate and faster than the SIR. We showed that a grid with 50 to 60 nodes leads to an MAPE of 0.1%, on average, for the S&P 500 Index—a threshold that is reasonable. As an application of the method, we also found the maximum likelihood estimates of the SVCJSI model and its nested models with the S&P 500 daily returns. The parameters obtained are consistent with those estimated in the literature.

Finally, this method is not only helpful for maximum likelihood estimation, but also for Bayesian inference as the likelihood function is an important building block when calculating posterior distributions. For instance, it would be interesting to combine the DNF with Markov chain Monte Carlo (MCMC) methods; this is left for future research.



**Figure 8: Filtered volatility, jump intensity, return jumps and variance jumps for the S&P 500 Index.** This figure reports the filtered volatility (i.e.,  $\sqrt{V_t/252}$ ), the jump intensity, the return jumps and the variance jumps for the S&P 500 Index and for the four models considered in this study. We use the S&P 500 return series from 1990 to 2018.

# Supplementary material

## A Particle filter

In this section, we briefly describe the sequential Monte Carlo method used to compute the likelihood function in this study. This implementation of the particle filter closely follows the work of Gordon et al. (1993); specifically, we use the bootstrap filter. Algorithm SM1 describes the various steps of the general SIR method.

---

### Algorithm SM1 Sequential importance resampling

---

- 1: initiate the state,  $\mathbf{x}_0^{(i)} \sim p(\cdot)$  for  $i \in \{1, 2, \dots, N_p\}$ , where  $p$  is the initial state density
  - 2: **for** each  $t \in \{1, \dots, T\}$  **do**
  - 3:   **for** each  $i \in \{1, \dots, N_p\}$  **do**
  - 4:     sample  $\tilde{\mathbf{x}}_t^{(i)} \sim q(\cdot \mid \mathbf{x}_{t-1}^{(i)})$
  - 5:     set  $\tilde{\mathbf{x}}_{0:t}^{(i)} = (\mathbf{x}_{0:t-1}^{(i)}, \tilde{\mathbf{x}}_t^{(i)})$
  - 6:     evaluate the importance weights, i.e.,  $w_t(\tilde{\mathbf{x}}_{0:t}^{(i)}) = w_{t-1}^{(i)} r(\mathbf{y}_t \mid \tilde{\mathbf{x}}_t^{(i)})$
  - 7:     normalize the importance weights, i.e.,  $\hat{w}_t^{(i)} = \frac{w_t(\tilde{\mathbf{x}}_{0:t}^{(i)})}{\sum_{j=1}^{N_p} w_t(\tilde{\mathbf{x}}_{0:t}^{(j)})}$
  - 8:   **end for**
  - 9:   resample with replacement  $N_p$  particles from  $\tilde{\mathbf{x}}_{0:t}^{(i)}$  according to the weights  $\hat{w}_t^{(i)}$
  - 10:   set  $w_t^{(i)} = \frac{1}{N_p}$  for each  $i \in \{1, \dots, N_p\}$
  - 11:   compute the time- $t$  likelihood contribution, i.e.,  $f(\mathbf{y}_t \mid \mathbf{y}_{1:t-1}) \approx \sum_{i=1}^{N_p} w_t(\tilde{\mathbf{x}}_{0:t}^{(i)})$
  - 12: **end for**
  - 13: compute the likelihood function  $\mathcal{L}(\Theta)$  by taking the product of the likelihood contributions
- 

## B Study of the DNF bias

The efficacy of the DNF method for estimating the parameters of jump-diffusion models using maximum likelihood estimation is now assessed. In this experiment, we compare the parameter estimates obtained by using DNF with their respective *true* values. The primary objective of this simulation exercise is to assess whether the DNF leads to systematic biases in the estimation parameters.

Specifically, for a set of likely parameters—consistent with most studies cited in the introduction of this study—we generate 100 ten-year return paths, i.e.,  $T = 2,520$ . Then, we use the DNF method coupled with a numerical optimizer to find the maximum likelihood estimates.

Table SM1 reports the average value, the bias as well as the root mean square error for each parameter and each model. In summary, the average is very close to the true value for most of the parameters, leading to small biases. The only exception is parameter  $\kappa$ : the bias associated with this parameter hovers between 0.25 and 0.75. Yet, even though this bias looks large at first sight, it is not so dreadful. In fact, when expressed in terms of daily persistence of the process—as it is typically done in econometrics—the difference is rather small, e.g.,  $e^{-3.717h} = 0.9882$  versus  $e^{-3.000h} = 0.9854$  for the SV model.

Table SM1: Study of the bias per model.

Panel A: SV model.				
	Average	True Value	Bias	RMSE
$\mu$	0.033	0.060	-0.027	0.042
$\kappa$	3.717	3.000	0.717	1.117
$\theta$	0.033	0.030	0.003	0.005
$\sigma$	0.300	0.300	0.000	0.031
$\rho_v$	-0.599	-0.600	0.001	0.070
Panel B: SVYJ model.				
	Average	True Value	Bias	RMSE
$\mu$	0.040	0.060	-0.020	0.045
$\kappa$	3.693	3.000	0.693	1.116
$\theta$	0.033	0.030	0.003	0.005
$\sigma$	0.300	0.300	0.000	0.033
$\rho_v$	-0.599	-0.600	0.001	0.081
$\theta$	4.909	5.000	-0.091	1.780
$\alpha$	-0.022	-0.020	-0.002	0.009
$\delta$	0.030	0.030	0.000	0.005
Panel C: SVCJ model.				
	Average	True Value	Bias	RMSE
$\mu$	0.059	0.060	-0.001	0.038
$\kappa$	3.530	3.000	0.530	0.778
$\theta$	0.031	0.030	0.001	0.008
$\sigma$	0.305	0.300	0.005	0.043
$\rho_v$	-0.598	-0.600	0.002	0.105
$\omega$	5.122	5.000	0.122	1.791
$\alpha$	-0.024	-0.020	-0.004	0.012
$\delta$	0.026	0.030	-0.004	0.007
$\nu$	0.011	0.010	0.001	0.005
$\rho_z$	-0.887	-1.000	0.113	0.591
Panel D: SVCJSI model.				
	Average	True Value	Bias	RMSE
$\mu$	0.058	0.060	-0.002	0.010
$\kappa$	3.261	3.000	0.261	0.451
$\theta$	0.031	0.030	0.001	0.005
$\sigma$	0.276	0.300	-0.024	0.033
$\rho_v$	-0.577	-0.600	0.023	0.091
$\chi$	3.123	3.000	0.123	0.373
$\omega$	5.060	5.000	0.060	0.780
$\xi$	5.089	5.000	0.089	0.817
$\rho_l$	-0.309	-0.300	-0.009	0.037
$\alpha$	-0.020	-0.020	0.000	0.003
$\delta$	0.028	0.030	-0.002	0.005
$\nu$	0.010	0.010	0.000	0.002
$\rho_z$	-1.004	-1.000	-0.004	0.130

This table reports the average parameter, the true value, the bias as well as the root mean square error, denoted by RMSE. For a set of likely parameters, we generate 100 ten-year return paths, i.e.,  $T = 2,520$ . Then, we use the DNF method coupled with a numerical optimizer to find the maximum likelihood estimates.

## References

- Aït-Sahalia, Y. 2002. Maximum likelihood estimation of discretely sampled diffusions: A closed-form approximation approach. *Econometrica* 70:223–262.
- Aït-Sahalia, Y., and R. Kimmel. 2007. Maximum likelihood estimation of stochastic volatility models. *Journal of Financial Economics* 83:413–452.
- Andersen, T. G., L. Benzoni, and J. Lund. 2002. An empirical investigation of continuous-time equity return models. *Journal of Finance* 57:1239–1284.
- Andersen, T. G., and B. E. Sørensen. 1996. GMM estimation of a stochastic volatility model: A Monte Carlo study. *Journal of Business & Economic Statistics* 14:328–352.

- Bakshi, G., C. Cao, and Z. Chen. 1997. Empirical performance of alternative option pricing models. *Journal of Finance* 52:2003–2049.
- Bardgett, C., E. Gourier, and M. Leippold. 2019. Inferring volatility dynamics and risk premia from the S&P 500 and VIX markets. *Journal of Financial Economics* 131:593–618.
- Bartolucci, F., and G. De Luca. 2001. Maximum likelihood estimation of a latent variable time-series model. *Applied Stochastic Models in Business and Industry* 17:5–17.
- Bates, D. S. 1996. Jumps and stochastic volatility: Exchange rate processes implicit in Deutsche Mark options. *Review of Financial Studies* 9:69–107.
- Bates, D. S. 2000. Post-'87 crash fears in the S&P 500 futures option market. *Journal of Econometrics* 94:181–238.
- Bates, D. S. 2006. Maximum likelihood estimation of latent affine processes. *Review of Financial Studies* 19:909–965.
- Bégin, J.-F., C. Dorion, and G. Gauthier. 2019. Idiosyncratic jump risk matters: Evidence from equity returns and options. Forthcoming in the *Review of Financial Studies*.
- Brandt, M. W., and P. Santa-Clara. 2002. Simulated likelihood estimation of diffusions with an application to exchange rate dynamics in incomplete markets. *Journal of Financial Economics* 63:161–210.
- Christoffersen, P., K. Jacobs, and K. Mimouni. 2010. Volatility dynamics for the S&P 500: Evidence from realized volatility, daily returns, and option prices. *Review of Financial Studies* 23:3141–3189.
- Christoffersen, P., K. Jacobs, and C. Ornthanalai. 2012. Dynamic jump intensities and risk premiums: Evidence from S&P 500 returns and options. *Journal of Financial Economics* 106:447–472.
- Clements, A., S. Hurn, and S. White. 2004. Discretised non-linear filtering for dynamic latent variable models: With application to stochastic volatility. Working Paper.
- Clements, A., S. Hurn, S. White, et al. 2006a. Estimating stochastic volatility models using a discrete non-linear filter. Working Paper.
- Clements, A. E., S. Hurn, and S. I. White. 2006b. Mixture distribution-based forecasting using stochastic volatility models. *Applied Stochastic Models in Business and Industry* 22:547–557.
- Cont, R. 2007. Volatility clustering in financial markets: Empirical facts and agent-based models. In *Long Memory in Economics*, pp. 289–309. Springer.
- Creal, D. 2012. A survey of sequential Monte Carlo methods for economics and finance. *Econometric Reviews* 31:245–296.
- Danielsson, J. 1994. Stochastic volatility in asset prices estimation with simulated maximum likelihood. *Journal of Econometrics* 64:375–400.
- Danielsson, J., and J.-F. Richard. 1993. Accelerated Gaussian importance sampler with application to dynamic latent variable models. *Journal of Applied Econometrics* 8:S153–S173.
- Duffie, D., J. Pan, and K. Singleton. 2000. Transform analysis and asset pricing for affine jump-diffusions. *Econometrica* 68:1343–1376.
- Eraker, B. 2001. MCMC analysis of diffusion models with application to finance. *Journal of Business & Economic Statistics* 19:177–191.
- Eraker, B., M. Johannes, and N. Polson. 2003. The impact of jumps in volatility and returns. *Journal of Finance* 58:1269–1300.
- Fridman, M., and L. Harris. 1998. A maximum likelihood approach for non-Gaussian stochastic volatility models. *Journal of Business & Economic Statistics* 16:284–291.
- Gordon, N. J., D. J. Salmond, and A. F. Smith. 1993. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. In *IEEE Proceedings F (Radar and Signal Processing)*, 140:107–113.

- Harvey, A., E. Ruiz, and N. Shephard. 1994. Multivariate stochastic variance models. *Review of Economic Studies* 61:247-264.
- Heston, S. 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies* 6:327.
- Hurn, A. S., K. A. Lindsay, and A. J. McClelland. 2015. Estimating the parameters of stochastic volatility models using option price data. *Journal of Business & Economic Statistics* 33:579-594.
- Jacquier, E., N. G. Polson, and P. E. Rossi. 1994. Bayesian analysis of stochastic volatility models. *Journal of Business & Economic Statistics* 12:371-389.
- Johannes, M., R. Kumar, and N. G. Polson. 1999. State dependent jump models: How do US equity indices jump. Working paper.
- Johannes, M. S., N. G. Polson, and J. R. Stroud. 2009. Optimal filtering of jump diffusions: Extracting latent states from asset prices. *Review of Financial Studies* 22:2759-2799.
- Kalman, R. E. 1960. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering* 82:35-45.
- Kitagawa, G. 1987. Non-Gaussian state-space modeling of nonstationary time series. *Journal of the American Statistical Association* 82:1032-1041.
- Langrock, R., I. L. MacDonald, and W. Zucchini. 2012. Some nonstandard stochastic volatility models and their estimation using structured hidden Markov models. *Journal of Empirical Finance* 19:147-161.
- Maheu, J. M., and T. H. McCurdy. 2004. News arrival, jump dynamics, and volatility components for individual stock returns. *Journal of Finance* 59:755-793.
- Malik, S., and M. K. Pitt. 2011. Particle filters for continuous likelihood evaluation and maximisation. *Journal of Econometrics* 165:190-209.
- Melino, A., and S. M. Turnbull. 1990. Pricing foreign currency options with stochastic volatility. *Journal of Econometrics* 45:239-265.
- Merton, R. C. 1976. Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics* 3:125-144.
- Nelson, D. B. 1988. The time series behavior of stock market volatility and returns. Ph.D. thesis, Massachusetts Institute of Technology.
- Ornthanalai, C. 2014. Levy jump risk: Evidence from options and returns. *Journal of Financial Economics* 112:69-90.
- Pan, J. 2002. The jump-risk premia implicit in options: Evidence from an integrated time-series study. *Journal of Financial Economics* 63:3-50.
- Pitt, M. K., S. Malik, and A. Doucet. 2014. Simulated likelihood inference for stochastic volatility models using continuous particle filtering. *Annals of the Institute of Statistical Mathematics* 66:527-552.
- Rémillard, B. 2013. *Statistical methods for financial engineering*. Chapman & Hall.
- Shephard, N. 1993. Fitting nonlinear time-series models with applications to stochastic variance models. *Journal of Applied Econometrics* 8:S135-S152.
- Taylor, S. J. 1986. *Modelling financial time series*. World Scientific.
- Todorov, V., and G. Tauchen. 2011. Volatility jumps. *Journal of Business & Economic Statistics* 29:356-371.
- Watanabe, T. 1999. A non-linear filtering approach to stochastic volatility models with an application to daily stock returns. *Journal of Applied Econometrics* 14:101-121.