Price bias and common practice in option pricing

J.-F. Bégin,
G. Gauthier

G–2018–56

July 2018

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, 2018
– Bibliothèque et Archives Canada, 2018

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, 2018
– Library and Archives Canada, 2018

GERAD HEC Montréal
3000, chemin de la Côte-Sainte-Catherine
Montréal (Québec) Canada H3T 2A7

Tél. : 514 340-6053
Téléc. : 514 340-5665
info@gerad.ca
www.gerad.ca

Price bias and common practice in option pricing

Jean-François Bégin [a]

Geneviève Gauthier [b]

[a] *Department of Statistics and Actuarial Science, Simon Fraser University, Burnaby (British Columbia) Canada, V5A 1S6*

[b] *GERAD & Department of Decision Sciences, HEC Montréal, Montréal (Québec) Canada, H3T 2A7*

jbegin@sfu.ca
genevieve.gauthier@hec.ca

July 2018
Les Cahiers du GERAD
G–2018–56

Abstract: Generally, the semiclosed-form option pricing formula for complex financial models depends on unobservable factors such as stochastic volatility and jump intensity. A popular practice is to use an estimate of these latent factors to compute the option price. However, in many situations, this plug-and-play approximation does not yield the appropriate price. This paper examines this bias and quantifies its impacts. We decompose the bias into terms that are related to the bias on the unobservable factors and to the precision of their point estimators. The approximated price is found to be highly biased when only the history of the stock price is used to recover the latent states. This bias is corrected when option prices are added to the sample used to recover the states' best estimate. We also show numerically that such a bias is propagated on calibrated parameters, leading to erroneous values.

Keywords: Option pricing, jump-diffusions, model calibration, estimation bias, information set

Acknowledgments: For their useful comments, the authors would like to thank Diego Amaya, Mathieu Boudreault, Christian Dorion, Frédéric Godin and Bruno Rémillard. Research supported in part by NSERC (Bégin and Gauthier), HEC Montréal (Gauthier) and Simon Fraser University (Bégin).

1 Introduction

Nowadays, the most common option pricing models are based on complex latent modelling factors such as stochastic volatility and jump intensity. In his seminal paper, Heston (1993) proposes the square root process to model the instantaneous variance of an asset. One of the main advantages of his framework is that vanilla option prices can be computed by a Fourier inversion of the conditional characteristic function, known in closed form for the square root diffusion. Return jumps—also latent—were considered by Bates (1996) and Bakshi et al. (1997), among others. Duffie et al. (2000, DPS henceforth) propose a very general class of affine jump-diffusion models that includes as subcases the usual extensions encountered in the literature.[1] Some other studies have considered non-affine frameworks, e.g., Jones (2003) and Christoffersen et al. (2010).[2]

In most of these models, the option pricing formula depends on the latent factors. The latter are not observable directly, which complicates the option pricing procedure. Indeed, one cannot observe the instantaneous variance or the jump intensity; these random variables remain unknown, even given the current time's information set, i.e., the history of the asset prices up to now. This paper's objective is to study the potential bias resulting from replacing the latent factors by some estimate—either the expected value of the latent variables (i.e., the filtered value) or, in a like manner, a proxy (e.g., the VIX index in the case of volatility).

In most jump-diffusion models, the expected (filtered) value of the latent factors can be estimated through a numerical procedure. For instance, if the state-space representation is linear and Gaussian, the Kalman (1960) filter provides the posterior distribution of model latent variables. In more complicated and general frameworks, the state-space representation is typically nonlinear, preventing us from using Kalman-type recursions. In these cases, one could use Markov chain Monte Carlo (MCMC) or sequential Monte Carlo (SMC) methods. On the one hand, MCMC methods provide a Bayesian solution to the inference problem, yielding the distribution of latent variables, conditional on observed data. They were used by Jones (1998), Eraker (2001) and Eraker (2004), for instance. On the other hand, filtering methods have received much attention recently as they are extremely flexible (e.g., Johannes et al., 2009; Ornthanalai, 2014; Bardgett et al., 2016). These methods also allow for a simulation-based distribution of latent variables, conditional on the observations. In this paper, we employ an SMC method by the means of a particle filter—the bootstrap one—that allows us to extract the posterior distribution of the model's latent variables as well as the *a posteriori* expected value of the latent factors.

The information used to extract the best estimate (of the latent factor) is of paramount importance. The tradition of extracting latent factors using option prices is already well anchored in the option pricing literature. This tradition dates back to Bakshi et al. (1997), who treat the latent factors as additional parameters. Using a two-stage procedure, Bates (2000) maximizes a likelihood function based on options and estimate volatility paths; this allows for consistency across physical and risk-neutral parameters; Huang and Wu (2004), as well as Christoffersen et al. (2009), use similar approaches. A different methodology allowed Pan (2002) and Santa-Clara and Yan (2010) to recover *perfectly* the latent factors from option prices. Finally, other papers, like ours, have relied on filtering to incorporate option information.[3] Typically, in filtering schemes, it is necessary to assume a pricing error to circumvent the possibility of singularity.

In this paper, we show that in a general jump-diffusion framework, the option price bias is a function of the bias on the latent variables' estimators, as well as its precision. Then, in a numerical study, we use Heston's (1993) framework to quantify the option bias and discuss its impact. Under this model and when only stock prices are used to infer the instantaneous variance, we find that a negative variance volga—the second derivative of the option pricing formula with respect to the instantaneous variance—yields a positive bias, and vice versa. At-the-money (ATM) options tend to have negative volgas, whereas out-of-the-money (OTM) and in-the-money (ITM) options have positive volgas. Consistent with the theory, we have found that approximated ATM option prices are larger than their true counterparts, producing a positive bias for these options, and that OTM and ITM options generate negative biases. When options are included in

[1] Subcases of DPS were suggested by Eraker (2004) and Sepp (2008), among others.

[2] However, these non-affine models do not allow for option prices in semiclosed-form solutions and are therefore less convenient from a computational standpoint. They could, however, still be used, and they suffer from the issue we discuss in this very paper.

[3] Carr and Wu (2010), Johannes et al. (2009), Christoffersen et al. (2012), Ornthanalai (2014), and Bardgett et al. (2016) to name a few.

the calculation of the filtered instantaneous variance, the bias tends to disappear as the number of options increases, although a single option is not enough to eliminate the bias. Indeed, when only one ATM option is used, the option bias is negative and proportional to the variance vega of the option.

As a robustness test, the numerical study is broadened to allow for jumps, i.e., Bates's (1996) stochastic volatility with jumps (SVJ) model is used. All the results are qualitatively consistent with those obtained with the Heston model: when no options are used, ATM options are positively biased whereas ITM and OTM options are negatively biased. The use of a single option is not enough to mitigate the bias; the latter vanishes as the number of options increases.

Regarding the impact of the bias on parameter estimates, our results also show that, for both the Heston model and the SVJ model, the popular practice of using the latent factors' estimates have the potential to yield biased parameters if option prices are excluded from the information set.

Finally, we investigate the impact of using other proxies of instantaneous variance—namely, the short-term ATM implied volatility and the VIX index. When these proxies are used, the option prices are highly biased and, as expected, the latent states' estimate errors are propagated into the calibrated parameters.

This paper is organized as follows. Section 2 presents the framework and discusses the issue at hand. It also explains the relationship between the option pricing formula and the bias. Simulation-based illustrations are presented in Section 3, whereas empirical results are discussed in Section 4. Section 5 shows the impact of using other proxies. Section 6 concludes.

2 Common practice in option pricing

2.1 The framework

To capture the behaviour of the asset price dynamics more efficiently, many market models include latent variables. Under the physical measure $\mathbb{P}$, the asset price S_t and the vector of latent factors $\boldsymbol{X}_t$ of size n are generally modelled by a jump-diffusion process of the form

$$dS_t = \mu_t(S_t, \boldsymbol{X}_t)\,dt + \sigma_t(S_t, \boldsymbol{X}_t)\,d\boldsymbol{W}_t + dJ_{S,t}, \tag{1}$$

$$d\boldsymbol{X}_t = \xi_t(S_t, \boldsymbol{X}_t)\,dt + \delta_t(S_t, \boldsymbol{X}_t)\,d\boldsymbol{W}_t + d\boldsymbol{J}_{X,t}, \tag{2}$$

where the drift μ and ξ, as well as the diffusion terms σ and δ, are real-valued functions of the model components. Moreover, $\boldsymbol{W}$ is a $(k+1)-$dimensional Brownian motion and $\boldsymbol{J}_t = [\,J_{S,t} \quad \boldsymbol{J}_{X,t}^{\top}\,]^{\top}$ represents the vector of jump terms. Usually, these coefficients are chosen in such a way that the model is Markovian.

The information structure is crucial in understanding the bias in option prices. Specifically, two filtrations are of special interest in this paper. First, the model information structure is characterized by the filtration $\{\mathcal{M}_t\}_{t\geq 0}$, where $\mathcal{M}_t = \sigma\{\boldsymbol{W}_u, \boldsymbol{J}_u : 0 \leq u \leq t\}$ is the σ-field generated by the Brownian motions and the jumps until time t. Second, the observation filtration is constructed from the σ-field generated by the underlying process and the option prices, i.e., $\mathcal{F}_t^M = \sigma\left\{S_{t_i}, \{C_{t_i,j}^{\text{Market}}\}_{j=1}^M : 0 \leq t_1 \leq t_2 \leq ... \leq t_n \leq t\right\}$, where $t_1 \leq t_2 \leq ... \leq t_n$ are times when we observe the price of the financial assets. More precisely, $C_{t_i,j}^{\text{Market}}$ denotes the market value of the $j^{\text{ th}}$ option at time t_i and M is the number of options used at each time. For most models, and depending on the size of M, the observation filtration $\mathcal{F}_t^M$ is coarser than the model's.

Finally, since the market model is incomplete, there are infinitely many martingale measures. In the following, $\mathbb{Q}$ represents one of these measures.

2.2 Option pricing

In the jump-diffusion framework of Equations (1) and (2), the model pricing formula depends on the actual price S_t and the latent factors $\boldsymbol{X}_t$:

$$C_t(S_t, \boldsymbol{X}_t) = \mathbb{E}^{\mathbb{Q}}\left[e^{-r(T-t)}g(S_T)\,\middle|\,\mathcal{M}_t\right], \tag{3}$$

where $g(S_T)$ is the time T payoff. In the general continuous time affine model framework, the call option valuation formula is given by the following expression:

$$C_t(S_t, \boldsymbol{X}_t) = S_t e^{-q(T-t)} P_1(S_t, \boldsymbol{X}_t) - K e^{-r(T-t)} P_2(S_t, \boldsymbol{X}_t),$$

where K is the strike price, T denotes the maturity,

$$P_z(S_t, \boldsymbol{X}_t) = \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \mathrm{Re} \left(\frac{e^{-\mathrm{i}u \log(K) - \log(S_t)} \exp\left(\mathcal{A}(u\mathrm{i} + (z-1), T-t) + \mathcal{B}(u\mathrm{i} + (z-1), T-t)^\top \boldsymbol{X}_t\right)}{u\mathrm{i}} \right) du,$$

and both $\mathcal{A}(u, T-t)$ and $\mathcal{B}(u, T-t)$ are model-dependent factors (see Duffie et al., 2000, for more details on these factors, and Appendix A.1 for an example).

2.3 The option bias

The state variable $\boldsymbol{X}_t$ is unobserved and, unfortunately, required to obtain the option price of Equation (3). One common way of dealing with this caveat is to replace it with an estimate. One easy avenue is to simply replace $\boldsymbol{X}_t$ by some proxy (e.g., the VIX index in the case of the volatility). Yet, a more proper approach is to extract an estimate of the latent factors from today's information. In this article, the conditional expected value of $\boldsymbol{X}_t$ is used; it is obtained by means of a particle filter.[4] Therefore, this approximation of the option price is given by

$$C_t(S_t, \widehat{\boldsymbol{X}}_t), \tag{4}$$

where $\widehat{\boldsymbol{X}}_t \equiv \mathbb{E}^{\mathbb{P}}\left[\boldsymbol{X}_t \,\middle|\, \mathcal{F}_t^M\right]$ is the filtered value, meaning that, the *best estimate* depends explicitly on the information used in the filter. However, since the option pricing formula is not linear in $\boldsymbol{X}_t$, and since $\widehat{\boldsymbol{X}}_t$ might not equal $\boldsymbol{X}_t$, Equation (4) is only an approximation of the option price.

This paper's objective is to examine the potential bias resulting from replacing $\boldsymbol{X}_t$ with $\widehat{\boldsymbol{X}}_t$. For the remainder of this paper, the option price bias is defined as

$$\text{Option Bias}_t = \mathbb{E}^{\mathbb{P}}\left[C_t(S_t, \widehat{\boldsymbol{X}}_t) - C_t(S_t, \boldsymbol{X}_t) \,\middle|\, \mathcal{F}_t^M\right], \tag{5}$$

that is, the expected difference between the option price and its approximation under the physical probability measure $\mathbb{P}$ and according to the actual information $\mathcal{F}_t^M$. The expectation of Equation (5) is taken under the physical measure because we are interested in the actual bias, the one observed in the real world.

Using a Taylor expansion of $C_t(S_t, \boldsymbol{X}_t)$ about $\widehat{\boldsymbol{X}}_t$,[5,6]

$$C_t(S_t, \boldsymbol{X}_t) \cong C_t\left(S_t, \widehat{\boldsymbol{X}}_t\right) + \widehat{\boldsymbol{\nabla}}_t^\top \left(\boldsymbol{X}_t - \widehat{\boldsymbol{X}}_t\right) + \frac{1}{2}\left(\boldsymbol{X}_t - \widehat{\boldsymbol{X}}_t\right)^\top \widehat{\mathbf{H}}_t \left(\boldsymbol{X}_t - \widehat{\boldsymbol{X}}_t\right) \tag{6}$$

where the gradient

$$\widehat{\boldsymbol{\nabla}}_t = \left(\frac{\partial C_t}{\partial X^{(i)}}(S_t, \boldsymbol{X}) \bigg|_{\boldsymbol{X} = \widehat{\boldsymbol{X}}_t} \right)_{i \in \{1, \ldots, n\}}$$

is the column vector of the option price partial derivatives and the Hessian matrix

$$\widehat{\mathbf{H}}_t = \left(\frac{\partial^2 C_t}{\partial X^{(i)} \partial X^{(j)}}(S_t, \boldsymbol{X}) \bigg|_{\boldsymbol{X} = \widehat{\boldsymbol{X}}_t} \right)_{i,j \in \{1, \ldots, n\}}$$

[4] Note, however, that the use of another estimate is possible.

[5] Because the Taylor expansion of $C_t(S_t, \boldsymbol{X}_t)$ is developed around its point estimate, $\widehat{\boldsymbol{X}}_t$, it approximates the distribution of the option price with respect to the estimation error $\boldsymbol{\epsilon}_t = \widehat{\boldsymbol{X}}_t - \boldsymbol{X}_t$. This is a different exercise from explaining the relationship between the option price and the distribution of the underlying asset price S_t, which commonly exhibits skewness and kurtosis effects.

[6] Since it is preferable to use a convergent estimator, we use a second-order Taylor expansion that captures the bias and the precision of the estimator.

contains partial second derivatives evaluated at $(S_t, \widehat{\boldsymbol{X}}_t)$. Taking the expectation on both sides of Equation (6), we obtain an approximation of the option price bias associated with the use of a point estimate $\widehat{\boldsymbol{X}}_t$ of $\boldsymbol{X}_t$. Indeed, letting

$$\boldsymbol{\epsilon}_t = \widehat{\boldsymbol{X}}_t - \boldsymbol{X}_t$$

be the estimation error associated with the latent vector, the option price bias is

$$\begin{aligned}
\text{Option Bias}_t &= \mathbb{E}^{\mathbb{P}}\left[C_t\left(S_t, \widehat{\boldsymbol{X}}_t\right) - C_t\left(S_t, \boldsymbol{X}_t\right) \,\middle|\, \mathcal{F}_t^M \right] \\
&\cong \mathbb{E}^{\mathbb{P}}\left[\widehat{\boldsymbol{\nabla}}_t^\top \boldsymbol{\epsilon}_t \,\middle|\, \mathcal{F}_t^M \right] - \frac{1}{2}\mathbb{E}^{\mathbb{P}}\left[\boldsymbol{\epsilon}_t^\top \widehat{\mathbf{H}}_t^\top \boldsymbol{\epsilon}_t \,\middle|\, \mathcal{F}_t^M \right] \\
&= \sum_{i=1}^n \mathbb{E}^{\mathbb{P}}\left[\widehat{\nabla}_t^{(i)} \epsilon_t^{(i)} \,\middle|\, \mathcal{F}_t^M \right] - \frac{1}{2}\sum_{i=1}^n \sum_{j=1}^n \mathbb{E}^{\mathbb{P}}\left[\epsilon_t^{(i)} \, \widehat{\mathrm{H}}_t^{(i,j)} \epsilon_t^{(j)} \,\middle|\, \mathcal{F}_t^M \right].
\end{aligned}$$

If the point estimator $\widehat{\boldsymbol{X}}_t$ is $\mathcal{F}_t^M$-measurable, that is, it is known on the basis of the information contained in $\mathcal{F}_t^M$, then the option price bias becomes

$$\begin{aligned}
\text{Option Bias}_t &\cong \sum_{i=1}^n \widehat{\nabla}_t^{(i)} \mathbb{E}^{\mathbb{P}}\left[\epsilon_t^{(i)} \,\middle|\, \mathcal{F}_t^M \right] - \frac{1}{2}\sum_{i=1}^n \sum_{j=1}^n \widehat{\mathrm{H}}_t^{(i,j)} \mathbb{E}^{\mathbb{P}}\left[\epsilon_t^{(i)} \epsilon_t^{(j)} \,\middle|\, \mathcal{F}_t^M \right] \\
&= \underbrace{\sum_{i=1}^n \widehat{\nabla}_t^{(i)} \mathbb{E}^{\mathbb{P}}\left[\epsilon_t^{(i)} \,\middle|\, \mathcal{F}_t^M \right]}_{\text{First-Order Latent Variable Bias Term}} \underbrace{- \frac{1}{2}\sum_{i=1}^n \sum_{j=1}^n \widehat{\mathrm{H}}_t^{(i,j)} \mathbb{E}^{\mathbb{P}}\left[\epsilon_t^{(i)} \,\middle|\, \mathcal{F}_t^M \right] \mathbb{E}^{\mathbb{P}}\left[\epsilon_t^{(j)} \,\middle|\, \mathcal{F}_t^M \right]}_{\text{Second-Order Latent Variable Bias Term}} \\
&\quad \underbrace{- \frac{1}{2}\sum_{i=1}^n \sum_{j=1}^n \widehat{\mathrm{H}}_t^{(i,j)} \mathrm{Corr}^{\mathbb{P}}\left[\epsilon_t^{(i)}, \epsilon_t^{(j)} \,\middle|\, \mathcal{F}_t^M \right] \sqrt{\mathrm{Var}^{\mathbb{P}}\left[\epsilon_t^{(i)} \,\middle|\, \mathcal{F}_t^M \right]} \sqrt{\mathrm{Var}^{\mathbb{P}}\left[\epsilon_t^{(j)} \,\middle|\, \mathcal{F}_t^M \right]}}_{\text{Precision Term}}. \qquad (7)
\end{aligned}$$

The first two terms of the bias decomposition of Equation (7) are affected by the latent variable estimator biases, $\mathbb{E}^{\mathbb{P}}\left[\epsilon_t^{(i)} \,\middle|\, \mathcal{F}_t^M \right]$, whereas the third is linked to the precision of the latent variable point estimators, the standard deviation $\sqrt{\mathrm{Var}^{\mathbb{P}}\left[\epsilon_t^{(i)} \,\middle|\, \mathcal{F}_t^M \right]}$.

2.4 Bias in Heston's framework

To gain more intuition on the option price bias, we focus on Heston's (1993) framework in our analysis: this model is well known and one of the most widely used models by practitioners and academics as it allows for semiclosed-form solutions for option pricing.[7] Still, we feel compelled to stress that all the other jump-diffusions—or subcases of DPS, for that matter—suffer from the same issue and would yield similar biases.[8,9]

Under the physical measure $\mathbb{P}$, the asset price S_t and the latent instantaneous stochastic volatility $\boldsymbol{X}_t \equiv V_t$ satisfy:

$$dS_t = \left(r - q + \eta V_t\right) S_t \, dt + \sqrt{V_t} \, S_t \, dW_{S,t}, \qquad (8)$$

$$dV_t = \kappa\left(\theta - V_t\right) dt + \sigma\sqrt{V_t} \, dW_{V,t}, \qquad (9)$$

where W_S and W_V are two correlated Brownian motions such that $\langle W_S, W_V \rangle_t = \rho t$. Moreover, q is the instantaneous dividend yield, κ is the speed of mean reversion, θ is the long-run variance level, σ is the volatility parameter of the variance process, and η is the equity risk premium parameter.

[7] Using the moment generating function of $\log(S_T)$, one can recover the price of European vanilla options in semiclosed forms. Details on the option pricing formula are given in Appendix A.1.

[8] Appendix A.5 assesses the size of the bias when using the stochastic volatility with jumps (SVJ) model—as proposed by Bates (1996). For this model, the results are qualitatively similar to those obtained with the Heston model.

[9] In general, parameters are more precisely estimated as the length of the observation series increases; this remains true in the presence of latent factors. However, the estimation of latent variables is different because their precision is related to the noise-to-signal ratio. In that perspective, it is expected that having fewer latent factors, or more observations at each period, decreases the noise-to-signal ratio, that is, increases the precision of the filtered estimates and, therefore, reduces the bias in Equation (7). Although Equation (7) shows theoretically that the bias exists in general models when the filtered value is used in lieu of the latent factor, we study empirically the size of this bias in a simple framework where the noise-to-signal ratio should be smaller.

Under the pricing measure $\mathbb{Q}$, the dynamics are given as follows:

$$dS_t = (r - q)\, S_t\, dt + \sqrt{V_t} S_t\, dW_{S,t}^{\mathbb{Q}},$$
$$dV_t = \kappa^* \left(\theta^* - V_t\right) dt + \sigma \sqrt{V_t}\, dW_{V,t}^{\mathbb{Q}},$$

where the correlation coefficient is still given by ρ, i.e., $\langle W_S^{\mathbb{Q}}, W_V^{\mathbb{Q}} \rangle_t = \rho t$, $\kappa^* = \kappa + \lambda\sigma$, $\theta^* = \kappa\theta/\kappa^*$, and the volatility risk premium parameter is λ.

In Heston's framework, there is a single latent factor, the instantaneous stochastic variance V_t. The option bias decomposition of Equation (7) therefore reduces to

$$\begin{aligned} \text{Option Bias}_t \cong & \underbrace{\left(\frac{\partial C_t}{\partial V}(S_t, V) \bigg|_{V=\widehat{V}_t} \right) \mathbb{E}^{\mathbb{P}} \left[\epsilon_t \,\middle|\, \mathcal{F}_t^M \right]}_{\text{First-Order Latent Variable Bias Term}} \underbrace{- \frac{1}{2} \left(\frac{\partial^2 C_t}{\partial V^2}(S_t, V) \bigg|_{V=\widehat{V}_t} \right) \left(\mathbb{E}^{\mathbb{P}} \left[\epsilon_t \,\middle|\, \mathcal{F}_t^M \right] \right)^2}_{\text{Second-Order Latent Variable Bias Term}} \\ & \underbrace{- \frac{1}{2} \left(\frac{\partial^2 C_t}{\partial V^2}(S_t, V) \bigg|_{V=\widehat{V}_t} \right) \text{Var}^{\mathbb{P}} \left[\epsilon_t \,\middle|\, \mathcal{F}_t^M \right]}_{\text{Precision Term}}, \end{aligned} \tag{10}$$

where the estimation error $\epsilon_t = \widehat{V}_t - V_t$ is the difference between the estimated and true instantaneous variances.

2.4.1 First-order latent variable bias term

In Heston's framework, the gradient $\widehat{\nabla}_t = \frac{\partial C_t}{\partial V}(S_t, V)\big|_{V=\widehat{V}_t}$ is simply the variance vega, the first derivative of the option price with respect to the instantaneous variance.[10] To gain some insight into the sign of the first term of Equation (10), let us consider a simpler model. If we assume that the volatility of the variance parameter σ is nil, i.e., $\sigma = 0$, and that the long-run variance level θ is equal to V_t, then our framework falls back to the well-known Black and Scholes (1973, BS henceforth) model and $V_s = V$, a constant for all s, is the squared volatility (or variance). Under the BS model, the value of the variance vega is well known:

$$\text{Variance Vega} = e^{-q(T-t)} S_t \phi(d_1) \frac{\sqrt{T-t}}{2\sqrt{V}}, \tag{11}$$

$$d_1 = \frac{\log(S_t/K) + \left(r - q + \frac{1}{2}V\right)(T-t)}{\sqrt{V(T-t)}}, \tag{12}$$

where K is the strike price, T is the maturity and function $\phi(\cdot)$ is the probability density function of the standardized Gaussian distribution.

This first derivative is always greater than zero, meaning that the sign of the first term in the option bias should be the same as that of the latent variable bias $\mathbb{E}^{\mathbb{P}}\left[\epsilon_t \,\middle|\, \mathcal{F}_t^M\right]$. In addition, the bias on the approximated price should be larger for ATM options as $\phi(d_1)$ is worth more when $\log(S_t/K) + \left(r - q + \frac{1}{2}V\right)(T-t)$ is in the neighbourhood of zero. On the other hand, as the option becomes in- or out-of-the-money, its variance vega should decrease.

This behaviour remains true if $\sigma > 0$, i.e., for the Heston model, although it is virtually impossible to show analytically since the vega is only known in semiclosed form (see Appendix A.3). To compute the numerical value for Heston's vega, we use the parameters estimated by Aït-Sahalia and Kimmel (2007) as presented in Table 1.[11] We also assume that the risk-free rate r and the dividend yield q are zero, $V_t = 0.02$, and $S_t = 1000$.

[10] We use the term *variance vega* because the vega is the second derivative with respect to the volatility parameter, whereas our vega is calculated with respect to the instantaneous variance.

[11] The parameters taken from Aït-Sahalia and Kimmel (2007) are estimated using a maximum likelihood-based method, using only one option over a different sample period (i.e., 1990 to 2003). The parameter estimation via likelihood approaches is beyond the scope of this paper; rather, the latter focuses on the potential pricing bias when the uncertainty surrounding the estimation of the latent variables is omitted. Moreover, these parameters are consistent with those estimated by other studies (see Table 11 of Hurn et al., 2015, for an exhaustive review of these parameters).

Table 1: Estimated parameters using the Heston model on the S&P 500 over 1990–2004.

Parameter	η	λ	κ	θ	σ	ρ
Value	2.562	0.000	5.130	0.044	0.520	-0.754

The parameters in this table are those estimated by Aït-Sahalia and Kimmel (2007). These parameters are estimated using a maximum likelihood-based method, using only one *option* over a different sample period (i.e., 1990 to 2003). One exception is λ: in Aït-Sahalia and Kimmel (2007), they do not estimate this parameter because it is too cumbersome to estimate in their setting. In fact, "the market price of risk parameter [λ] is held fixed, so that the value of [κ] (i.e., the $\mathbb{P}-$measure speed of mean reversion) is constrained by the value of [κ^*] (i.e., the $\mathbb{Q}-$measure speed of mean reversion)." Section SM.A.2 of the Supplementary Material assesses the impact on our results of using a different value.

The top panel of Figure 1 shows the vegas for the Heston model as a function of moneyness and time to maturity.[12] The behaviour of this surface is very similar to the one obtained under the BS model: the variance vega is always positive. In addition, ATM options have a larger variance vega, whereas ITM and OTM options tend to have a smaller vega. The first-order latent variable bias term should, therefore, have the same sign as the latent variable bias, $\mathbb{E}^{\mathbb{P}}\left[\epsilon_t \,\middle|\, \mathcal{F}_t^M\right]$.

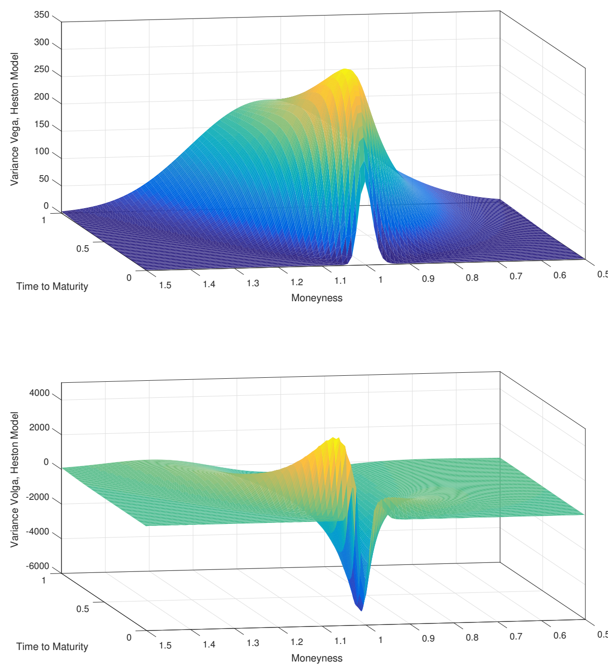

Figure 1: Variance vega and volga as a function of the time to maturity and the moneyness of the option for the Heston model. The figure shows the first (top panel) and the second derivative (bottom panel) of the option price with respect to V, the variance under the Heston model. The moneyness is defined as the strike price divided by the forward price. The time to maturity is given in years.

2.4.2 Second-order latent variable bias term

According to Equation (10), the option bias depends on the second derivative of the option pricing formula, with respect to the instantaneous variance, the unidimensional Hessian, $\widehat{\mathrm{H}}_t = \frac{\partial^2 C_t}{\partial V^2}(S_t, V)\Big|_{V=\widehat{V}_t}$. To gain some insight, let us again assume that $\sigma = 0$ and $V_t = \theta$. In that case, it is easy to get the value of the variance volga, the second derivative of the option price with respect to the instantaneous variance.

Proposition 1 *Variance Volga Under the BS Model. The variance volga is given by*

$$\text{Variance Volga} = \frac{e^{-q(T-t)} S_t \sqrt{T-t}\,\phi(d_1)}{4V^{3/2}} \left(d_1^2 - d_1\sqrt{V(T-t)} - 1\right), \tag{13}$$

where d_1 is given in Equation (12).

[12] Throughout this paper, moneyness is defined as the strike price divided by the forward price, i.e., $Ke^{-r(T-t)}/S_t e^{-q(T-t)}$.

Proof. See Appendix A.2. □

From Equation (13), we see that the second derivative with respect to V changes sign twice; we can also find the zeros of Equation (13) to identify the moneyness value at which the volga changes sign for a given time to maturity. Indeed, the volga is negative when the moneyness is in the following interval:

$$\left[e^{(r-q)(T-t)-\frac{1}{2}\sqrt{V(T-t)}\sqrt{4+V(T-t)}},\ e^{(r-q)(T-t)+\frac{1}{2}\sqrt{V(T-t)}\sqrt{4+V(T-t)}}\right]. \tag{14}$$

The second derivative is therefore less than zero for ATM options.

Again, this behaviour remains true for the Heston model. The bottom panel of Figure 1 shows an example of volgas for the Heston model: ATM options have a negative variance volga, whereas ITM and OTM options tend to have a positive variance volga. Different parameters tend to yield similar patterns, except for correlation parameter ρ. Interestingly, in Figure 1, the volga's behaviour is asymmetric: the second derivative is much more positive for OTM calls (ITM puts) than ITM calls (OTM puts). This is a clear consequence of correlation parameter ρ. By using a correlation parameter of zero, the surface becomes symmetric; a positive ρ yields an asymmetric surface, where the volga of ITM calls (OTM puts) is larger than OTM calls (ITM puts).[13]

Given the numerical evaluation of the variance volga, we can find the regions where the volga is positive and negative, similar to the region defined in (14). However, this now requires numerical inversion of a function defined through numerical integration. Figure 2 exhibits the regions where the variance volga is negative and positive for the Heston model (and for the parameters of Table 1). These regions are also very similar to those obtained with the BS model (i.e., Equation (14)), further supporting the insight developed at the beginning of this subsection.

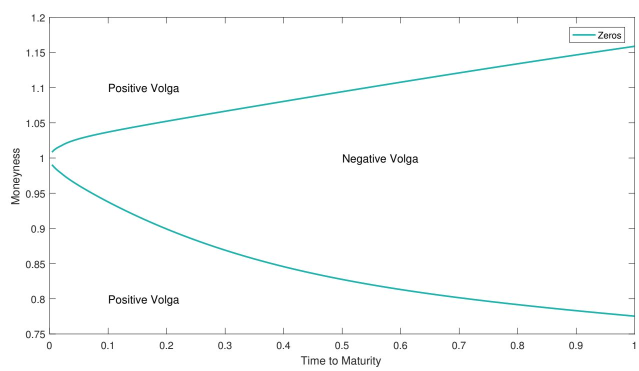

Figure 2: Regions where the variance volga is positive and negative for the Heston model. The figure shows the regions where the second derivative of the option price with respect to V is positive and negative.

2.4.3 Precision term

The third term of Equation (10) is related to the precision of the estimate, $\sqrt{\mathrm{Var}^{\mathbb{P}}\left[\epsilon_t \,\middle|\, \mathcal{F}_t^M\right]}$, and to the variance volga of the option. Given that the variance estimate is always positive, then this third effect and the variance volga differ in their sign, and the option bias surface is expected to be proportional to the inverse of the option volga.

Summing up the effects, all three terms contribute to the option price bias when an observable proxy is used. Indeed, replacing the instantaneous variance with the squared VIX or the squared implied volatility of a given option for example, yields $\mathbb{E}^{\mathbb{P}}\left[\epsilon_t \,\middle|\, \mathcal{F}_t^M\right] \neq 0$ and, therefore, all three terms contribute to the option price bias.

[13] The figures are available on request.

On the other hand, when the instantaneous variance is estimated from the filtered variance, that is $\widehat{V}_t = \mathbb{E}^{\mathbb{P}}\left[V_t \,\middle|\, \mathcal{F}_t^M\right]$, then

$$\mathbb{E}^{\mathbb{P}}\left[\epsilon_t \,\middle|\, \mathcal{F}_t^M\right] = \mathbb{E}^{\mathbb{P}}\left[\widehat{V}_t - V_t \,\middle|\, \mathcal{F}_t^M\right] = \mathbb{E}^{\mathbb{P}}\left[\mathbb{E}^{\mathbb{P}}\left[V_t \,\middle|\, \mathcal{F}_t^M\right] - V_t \,\middle|\, \mathcal{F}_t^M\right] = 0,$$

in theory. Therefore, the option price bias should only depend on the precision of the instantaneous variance estimator. Considering the multiplicative factor in front of the third term of Equation (10), the option bias surface is expected to be proportional to the inverse of the option volga.

However, in practice, the picture is less clear-cut: there might be a disagreement between the observed option prices and the model ones, that is, the observed option price is a function h of the model price and a noise term ν_t:

$$C_t^{\text{Market}} = h\left(C_t\left(S_t, V_t\right), \nu_t\right).$$

If the instantaneous variance is filtered using a single option, then the presence of the noise term ν_t creates a distortion that makes $\mathbb{E}^{\mathbb{P}}\left[\epsilon_t \,\middle|\, \mathcal{F}_t^M\right] \neq 0$. If the option noise term is dominant, then the option price bias of Equation (10) is dominated by the first-order latent variable bias term and the option price bias surface will be proportional to the option variance vega. On the one hand, when several options are used in the filtering process, the noise term effect is mitigated, both the bias and the precision of the instantaneous variance estimator are improved and the option price bias vanishes.[14] On the other hand, when no options are used, then we expect the precision to be rather poor, and therefore the option price bias to be driven by the precision term, e.g., the variance volga.

3 Simulation experiments

3.1 Option price bias

We have shown that approximated prices have the potential of either being larger or smaller than their model values, depending on the moneyness of the option, on the size of the estimate's precision and on the importance of the latent instantaneous variance bias. This section now investigates the impact of using the approximation instead of the price of Equation (3); this impact is quantified by calculating the option price bias.

To compute the approximated price of Equation (4), we need the expected value of the instantaneous variance, given the history of underlying and option prices. In this paper, the bootstrap filter of Gordon et al. (1993) (see Appendix A.4) provides a simulated approximation of the posterior density $f_{\mathbb{P}}\left(V_t \,\middle|\, \mathcal{F}_t^M\right)$ for all t. It is, therefore, possible to obtain the filtered value of V_t, $\mathbb{E}^{\mathbb{P}}\left[V_t \,\middle|\, \mathcal{F}_t^M\right]$.

To recover a surface of the option biases, we propose a simulation-based exercise. First, one hundred daily paths of the underlying asset over one year (250 business days) are generated using the Aït-Sahalia and Kimmel (2007) parameters of Table 1 along with $r = 0$, $q = 0$, $V_t = 0.02$, and $S_t = 1000$. For each path, we compute the daily option prices to be included in the filtering scheme. We consider four different filtrations: (1) $\mathcal{F}_t^0$, a filtration that does not include option prices; (2) $\mathcal{F}_t^1$, a filtration that includes one option price on each day (ATM option with 30 days to maturity); (3) $\mathcal{F}_t^5$, a filtration that includes five option prices on each day (options with 30 days to maturity and moneyness of 0.80, 0.90, 1.00, 1.10 and 1.20); and (4) $\mathcal{F}_t^{10}$, a filtration that includes ten option prices on each day (options with 30 and 90 days to maturity, and moneyness of 0.80, 0.90, 1.00, 1.10 and 1.20). Table 2 summarizes the various filtrations used in this article.[15]

We then add pricing errors to the option prices. More specifically, we use the Error Specification 4 (ES4) of Hurn et al. (2015), a multiplicative pricing error specification centred at one, with a standard deviation

[14] Indeed, this is true if the noise terms are uncorrelated.

[15] In these filtrations, we focus on options with shorter maturities as they tend to yield better results. Results reported in Section SM.A.1 of the Supplementary Material show that using an option with a maturity of 150 days in the filter gives an absolute bias of about -0.25 for ATM options, which is about five times more than what we obtain if we use an option with a maturity of 30 days.

of α, such that[16]

$$C_t^{\text{Market}} = C_t(S_t, V_t)e^{-\frac{1}{2}\alpha^2+\alpha\nu_t}, \quad \nu_t \sim \mathcal{N}(0,1),$$

where C_t^{Market} is the *observed* market price.[17] This error specification is a "sensible choice because the mean value of the option price is not distorted by the choice of pricing error distribution."

Table 2: The information structure used in the simulation experiments and in the empirical investigation.

Name	Symbol	Description
Zero Options	$\mathcal{F}_t^0$	A filtration that includes zero option prices.
One Option	$\mathcal{F}_t^1$	A filtration that includes one option price on each day. This option has a maturity of 30 days and a moneyness of 1.00.
Five Options	$\mathcal{F}_t^5$	A filtration that includes five option prices on each day. These options have a maturity of 30 days and a moneyness of 0.80, 0.90, 1.00, 1.10 and 1.20.
Ten Options	$\mathcal{F}_t^{10}$	A filtration that includes ten option prices on each day. These options have a maturity of 30 and 90 days, and a moneyness of 0.80, 0.90, 1.00, 1.10 and 1.20.

For each path and each filtration, we apply the filtering methodology of Appendix A.4 and recover the time series of $\widehat{V}_t$. Finally, we compute option prices according to Equations (4) using filtered $\widehat{V}_t$ for different values of moneyness and time to maturity. Specifically, to recover a cross-section of biases, we calculate 5,050 option prices, each with a different moneyness and maturity, on each day: the moneyness buckets range from 0.50 to 1.50, with steps of 0.01, and the maturity buckets range from 5 days to 1 year, with steps of 5 days.

In the rest of this paper, we refer to two different types of errors: absolute and relative. The absolute error is simply given by

$$C_t(S_t, \widehat{V}_t) - C_t(S_t, V_t). \tag{15}$$

As a consequence of the put-call parity relationship, the absolute error on a call option and on a put option is the same:

$$\begin{aligned} C_t^{\text{Call}}(S_t, \widehat{V}_t) - C_t^{\text{Put}}(S_t, \widehat{V}_t) &= S_t e^{-q(T-t)} - Ke^{-r(T-t)}, \\ C_t^{\text{Call}}(S_t, V_t) - C_t^{\text{Put}}(S_t, V_t) &= S_t e^{-q(T-t)} - Ke^{-r(T-t)}, \\ \Rightarrow C_t^{\text{Call}}(S_t, \widehat{V}_t) - C_t^{\text{Call}}(S_t, V_t) &= C_t^{\text{Put}}(S_t, \widehat{V}_t) - C_t^{\text{Put}}(S_t, V_t). \end{aligned}$$

The relative error is defined as the absolute error divided by the option's model price:

$$\frac{C_t(S_t, \widehat{V}_t) - C_t(S_t, V_t)}{C_t(S_t, V_t)}. \tag{16}$$

The relative error now depends on the option type as the price of a call option is not equal to the price of a put option. We report average errors across the one hundred paths for given values of moneyness and time to maturity.

Figure 3 shows the absolute bias as a function of the maturity of the option and its moneyness for four different instantaneous variance estimates: using zero options (top left panel), one option (top right panel), five options (bottom left panel), and ten options (bottom right panel).[18]

When no options are used, the bias is positive for ATM options and is larger for options that have shorter maturities. The bias is negative for OTM and ITM options, again, displaying an asymmetric behaviour similar to that of the variance volga. Indeed, the top left panel of Figure 3 closely resembles to the bottom panel of Figure 1, except that the sign of the bias is the opposite of that of the variance volga. This observation is consistent with a large filtering variance on V_t, c.f. a large third term in Equation (10). When only one option is used on each day, the option bias is negative for each moneyness and time to maturity. The top right panel

[16] Parameter α is set to 0.05 in this section. This value is consistent with the one used in the simulation experiments done by Hurn et al. (2015).

[17] In this exercise, the observed market price corresponds to the simulated model price to which we add the pricing error. Below, in the empirical investigation of Section 4, it will be the actual price that is observed in the market.

[18] Note that these biases are not calculated based on the approximation of Equation (10). Instead, they are obtained by Monte Carlo simulation and therefore are not impacted by the truncation error associated with the Taylor expansion.

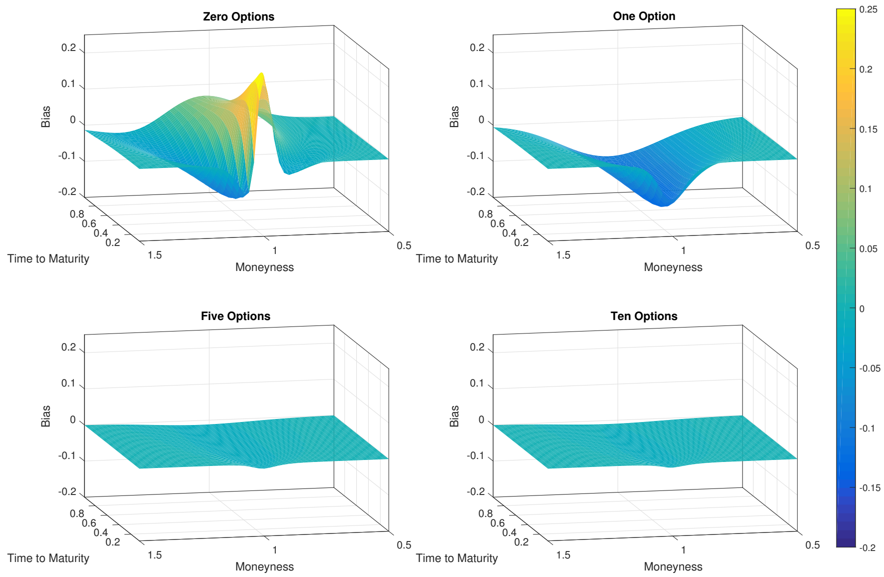

Figure 3: Option absolute bias as a function of the number of options used in the filter, the time to maturity and the moneyness of the option for the Heston model. This figure shows the absolute bias for options, as a function of the number of options used in the filter (i.e., zero, one, five and ten). For each simulated path and each filtration, we apply the filtering methodology to extract the best estimate (i.e., filtered value). We then compute the option price for various times to maturity and moneyness values using the approximation formula of Equations (4).

of Figure 3 is similar to the inverse of the variance vega as shown in the top panel of Figure 1, meaning that the second order influence is (almost) completely switched off when an option is added to the information set. This case is consistent with a first-order latent variable bias effect, where $\mathbb{E}^{\mathbb{P}}\left[\epsilon_t \,\middle|\, \mathcal{F}_t^1\right] < 0$. Lastly, when considering five and ten options, the option bias is virtually zero, which is consistent with Equation (10).

Tables 3 and 4 exhibit the relative biases for call and put options, stated as a percentage. The biases are averaged across moneyness and maturity buckets. Without options (Panel A), ATM call and put prices' biases are mildly positive, ranging from 0.1% to 1.0% on a relative basis. OTM call options tend to have a large negative relative bias that could reach about -32%, whereas ITM call options' relative bias is modest. On the other hand, OTM put options also have a large negative relative bias while ITM put options have a smaller negative relative bias. This result does not come as a surprise given that ITM options tend to have larger premiums than OTM options do, meaning that relative biases shrink as the denominator of Equation (16) becomes larger. When one option is included in the filter on each day (Panel B), the relative biases are negative for both call and put options. Again, OTM options tend to have large biases (in absolute size) that could reach -8%, yet these tend to be smaller than those obtained without any options. Finally, the cases using five and ten options are similar (Panels C and D) and the relative biases are modest for all the moneyness values and times to maturity considered.

Overall, without options, ATM options are positively impacted on an absolute basis, whereas OTM short-end options tend to have a large negative bias, on a relative basis. In the latter case, a small error of a few pennies might have dire consequences on the final pricing, relatively speaking. When options are included, the bias is negative and tends to become smaller as the number of options increases.

These conclusions still hold true if we consider the SVJ model: the bias is positive for ATM options and negative for OTM and ITM options when no options are used in the particle filter. As we increase the number of options in the filtration (i.e., $\mathcal{F}_t^1$, $\mathcal{F}_t^5$, and $\mathcal{F}_t^{10}$), the bias becomes negative and tends to disappear, consistent with a first-order latent variable bias effect observed with the Heston model.[19] The relative biases (shown in the Supplementary Material) are similar to those of Tables 3 and 4: OTM options tend to have the largest relative biases, as expected. Yet they tend to be less significant than those obtained in the Heston case—option prices, especially with short maturities, are larger under the SVJ model making the denominator of Equation (16) larger, and the relative bias smaller.

[19]See Appendix A.5.1 and Figure 7 for more details.

Table 3: Relative biases for calls as a function of time to maturity and moneyness for the Heston model.

Panel A: Zero Options.					
Moneyness \ Maturity	$[0.0, 0.2]$	$[0.2, 0.4]$	$[0.4, 0.6]$	$[0.6, 0.8]$	$[0.8, 1.0]$
$[0.650, 0.750]$	-0.001	-0.003	-0.002	-0.000	0.001
$[0.750, 0.850]$	-0.008	-0.003	0.005	0.008	0.009
$[0.850, 0.950]$	0.008	0.049	0.049	0.040	0.033
$[0.950, 0.975]$	0.235	0.201	0.130	0.089	0.066
$[0.975, 1.000]$	0.554	0.311	0.178	0.116	0.083
$[1.000, 1.025]$	0.978	0.439	0.231	0.145	0.101
$[1.025, 1.050]$	1.051	0.556	0.286	0.175	0.120
$[1.050, 1.150]$	-1.631	0.135	0.273	0.208	0.152
$[1.150, 1.250]$	-15.972	-5.284	-1.116	-0.130	0.074
$[1.250, 1.350]$	-32.313	-17.242	-6.595	-2.082	-0.610
Panel B: One Option.					
Moneyness \ Maturity	$[0.0, 0.2]$	$[0.2, 0.4]$	$[0.4, 0.6]$	$[0.6, 0.8]$	$[0.8, 1.0]$
$[0.650, 0.750]$	-0.001	-0.003	-0.006	-0.007	-0.008
$[0.750, 0.850]$	-0.006	-0.016	-0.021	-0.022	-0.022
$[0.850, 0.950]$	-0.045	-0.067	-0.066	-0.060	-0.053
$[0.950, 0.975]$	-0.176	-0.169	-0.137	-0.111	-0.092
$[0.975, 1.000]$	-0.306	-0.240	-0.179	-0.139	-0.113
$[1.000, 1.025]$	-0.525	-0.340	-0.234	-0.174	-0.137
$[1.025, 1.050]$	-0.899	-0.484	-0.306	-0.217	-0.166
$[1.050, 1.150]$	-1.813	-0.930	-0.529	-0.348	-0.252
$[1.150, 1.250]$	-4.800	-2.574	-1.327	-0.768	-0.503
$[1.250, 1.350]$	-8.302	-5.069	-2.793	-1.556	-0.948
Panel C: Five Options.					
Moneyness \ Maturity	$[0.0, 0.2]$	$[0.2, 0.4]$	$[0.4, 0.6]$	$[0.6, 0.8]$	$[0.8, 1.0]$
$[0.650, 0.750]$	-0.000	-0.000	-0.001	-0.001	-0.001
$[0.750, 0.850]$	-0.001	-0.002	-0.003	-0.003	-0.003
$[0.850, 0.950]$	-0.005	-0.008	-0.008	-0.007	-0.007
$[0.950, 0.975]$	-0.022	-0.021	-0.017	-0.014	-0.011
$[0.975, 1.000]$	-0.040	-0.030	-0.022	-0.017	-0.014
$[1.000, 1.025]$	-0.068	-0.043	-0.029	-0.021	-0.017
$[1.025, 1.050]$	-0.111	-0.060	-0.038	-0.027	-0.020
$[1.050, 1.150]$	-0.203	-0.110	-0.064	-0.042	-0.031
$[1.150, 1.250]$	-0.576	-0.289	-0.153	-0.091	-0.060
$[1.250, 1.350]$	-1.301	-0.650	-0.326	-0.181	-0.112
Panel D: Ten Options.					
Moneyness \ Maturity	$[0.0, 0.2]$	$[0.2, 0.4]$	$[0.4, 0.6]$	$[0.6, 0.8]$	$[0.8, 1.0]$
$[0.650, 0.750]$	-0.000	-0.000	-0.001	-0.001	-0.001
$[0.750, 0.850]$	-0.001	-0.002	-0.002	-0.002	-0.002
$[0.850, 0.950]$	-0.004	-0.007	-0.007	-0.006	-0.005
$[0.950, 0.975]$	-0.018	-0.017	-0.014	-0.011	-0.009
$[0.975, 1.000]$	-0.031	-0.024	-0.018	-0.014	-0.011
$[1.000, 1.025]$	-0.053	-0.034	-0.023	-0.017	-0.014
$[1.025, 1.050]$	-0.088	-0.048	-0.030	-0.022	-0.017
$[1.050, 1.150]$	-0.176	-0.091	-0.053	-0.035	-0.025
$[1.150, 1.250]$	-0.604	-0.276	-0.135	-0.077	-0.051
$[1.250, 1.350]$	-1.262	-0.667	-0.318	-0.164	-0.098

This table reports the average relative call option biases for each bin. The biases, calculated from Equation (16) for a large cross-section of options, are sorted into buckets of moneyness and maturity. The biases are reported as a percentage.

3.2 Impact on calibrated parameter λ

In the previous section, we discussed the magnitude of the option price bias. Yet, it is hard to assess the relevance of this bias. In this spirit, we propose a calibration exercise on simulated paths as the goal of this section is to show that the option price bias could be propagated in the calibrated parameter.[20] Specifically, we focus only on the volatility risk premium parameter λ in this exercise. We start by generating one hundred

[20] There is a way to obtain a likelihood function as a by-product of the particle filter. However, its construction and optimization is not trivial and beyond the scope of this paper. We rely on the popular calibration approach regularly used by practitioners.

Table 4: Relative biases for puts as a function of time to maturity and moneyness for the Heston model.

Panel A: Zero Options.

Moneyness \ Maturity	[0.0, 0.2]	[0.2, 0.4]	[0.4, 0.6]	[0.6, 0.8]	[0.8, 1.0]
[0.650, 0.750]	-7.065	-1.558	-0.288	-0.018	0.045
[0.750, 0.850]	-2.697	-0.183	0.147	0.149	0.122
[0.850, 0.950]	0.190	0.444	0.285	0.185	0.130
[0.950, 0.975]	0.943	0.496	0.263	0.163	0.112
[0.975, 1.000]	0.956	0.441	0.232	0.145	0.101
[1.000, 1.025]	0.671	0.347	0.193	0.124	0.088
[1.025, 1.050]	0.268	0.237	0.150	0.102	0.075
[1.050, 1.150]	-0.052	0.013	0.045	0.047	0.042
[1.150, 1.250]	-0.015	-0.032	-0.021	-0.005	0.005
[1.250, 1.350]	-0.001	-0.007	-0.012	-0.012	-0.008

Panel B: One Option.

Moneyness \ Maturity	[0.0, 0.2]	[0.2, 0.4]	[0.4, 0.6]	[0.6, 0.8]	[0.8, 1.0]
[0.650, 0.750]	-2.984	-1.514	-0.891	-0.594	-0.430
[0.750, 0.850]	-1.985	-1.023	-0.614	-0.418	-0.308
[0.850, 0.950]	-1.116	-0.613	-0.388	-0.273	-0.207
[0.950, 0.975]	-0.706	-0.416	-0.276	-0.201	-0.156
[0.975, 1.000]	-0.527	-0.340	-0.234	-0.175	-0.138
[1.000, 1.025]	-0.360	-0.268	-0.195	-0.149	-0.120
[1.025, 1.050]	-0.229	-0.206	-0.160	-0.127	-0.104
[1.050, 1.150]	-0.057	-0.089	-0.088	-0.079	-0.069
[1.150, 1.250]	-0.004	-0.016	-0.026	-0.030	-0.032
[1.250, 1.350]	-0.000	-0.002	-0.005	-0.009	-0.012

Panel C: Five Options.

Moneyness \ Maturity	[0.0, 0.2]	[0.2, 0.4]	[0.4, 0.6]	[0.6, 0.8]	[0.8, 1.0]
[0.650, 0.750]	-0.356	-0.178	-0.106	-0.071	-0.052
[0.750, 0.850]	-0.229	-0.121	-0.074	-0.051	-0.037
[0.850, 0.950]	-0.133	-0.075	-0.047	-0.033	-0.025
[0.950, 0.975]	-0.089	-0.052	-0.034	-0.025	-0.019
[0.975, 1.000]	-0.068	-0.043	-0.029	-0.021	-0.017
[1.000, 1.025]	-0.047	-0.034	-0.024	-0.018	-0.015
[1.025, 1.050]	-0.028	-0.026	-0.020	-0.016	-0.013
[1.050, 1.150]	-0.006	-0.011	-0.011	-0.010	-0.008
[1.150, 1.250]	-0.001	-0.002	-0.003	-0.004	-0.004
[1.250, 1.350]	-0.000	-0.000	-0.001	-0.001	-0.001

Panel D: Ten Options.

Moneyness \ Maturity	[0.0, 0.2]	[0.2, 0.4]	[0.4, 0.6]	[0.6, 0.8]	[0.8, 1.0]
[0.650, 0.750]	-0.343	-0.159	-0.091	-0.060	-0.044
[0.750, 0.850]	-0.208	-0.103	-0.062	-0.042	-0.031
[0.850, 0.950]	-0.111	-0.061	-0.039	-0.027	-0.021
[0.950, 0.975]	-0.071	-0.042	-0.028	-0.020	-0.016
[0.975, 1.000]	-0.054	-0.034	-0.023	-0.017	-0.014
[1.000, 1.025]	-0.037	-0.027	-0.019	-0.015	-0.012
[1.025, 1.050]	-0.022	-0.020	-0.016	-0.013	-0.010
[1.050, 1.150]	-0.006	-0.009	-0.009	-0.008	-0.007
[1.150, 1.250]	-0.001	-0.002	-0.003	-0.003	-0.003
[1.250, 1.350]	-0.000	-0.000	-0.001	-0.001	-0.001

This table reports the average relative put option biases for each bin. The risk-free rate r and the dividend yield q are assumed to zero. Moveover, $V_0 = 0.02$, and $S_0 = 1000$. For the other parameters, we use the ones estimated by Aït-Sahalia and Kimmel (2007). The moneyness is defined as the strike price divided by the forward price. The time to maturity is given in years. One hundred paths of one year are simulated for the underlying asset. For each path, we compute daily option prices to be included in the filtering scheme, consistently with the filtration exposed in Table 2. We also add pricing error in the option prices. Then, for each path, each day and each filtration, we apply the filtering methodology of Appendix A.4. The biases, calculated from Equation (16) for a large cross-section of options, are sorted into buckets of moneyness and maturity. The biases are stated as a percentage.

daily paths of the underlying asset, each with a length of one year. The option values needed for the four filtrations $\mathcal{F}_t^0$, $\mathcal{F}_t^1$, $\mathcal{F}_t^5$, and $\mathcal{F}_t^{10}$ are generated. We also generate a panel of daily options to be used in the calibration stage: for this panel, the moneyness buckets range from 0.80 to 1.20, with steps of 0.05, and the maturity buckets range from 30 days to 150 days, with steps of 30 days. In keeping with the empirical literature on option pricing, we focus only on ATM and OTM options. We also include error terms to the option prices, as discussed in the previous section. The scale parameter α of the error distribution is still 0.05.

We then calibrate the parameter λ for each path and each filtration. We first apply the filtering methodology to recover the filtered variance $\widehat{V}_t$ using the underlying and option price values (i.e., for each filtration, $\mathcal{F}_t^0$, $\mathcal{F}_t^1$, $\mathcal{F}_t^5$, and $\mathcal{F}_t^{10}$). Then, we compute

$$\chi^2 = \sum_{t \in \tau} \sum_{k=1}^{K} \left(\log\left(C_{t,k}^{\text{Market}}\right) - \log\left(C_{t,k}(S_t, \widehat{V}_t)\right)\right)^2, \tag{17}$$

where $\tau = \{\frac{1}{250}, ..., \frac{250}{250}\}$, $C_{t,k}^{\text{Market}}$ is the k^{th} simulated market option price on day t, $C_{t,k}(S_t, \widehat{V}_t)$ is the k^{th} approximated option price on day t, and K is the total number of options on a given day. The logarithm transformation of the option prices reduces the impact of ATM options in the minimization.[21,22] For more details on this minimization methodology, see Section 9.6 of Rebonato (2005). These two steps are performed in turn, repeatedly, until we find the $\widehat{\lambda}$ that minimizes χ^2. This is achieved via numerical optimization, by leaving parameter λ unrestrained.[23] Notice that, in this setting, the function χ^2 depends on λ, and indirectly on the filtered instantaneous variance estimates, which also depends on λ. Using a different λ parameter leads to different option prices, which in turn leads to a different estimate of V_t when we apply our filtering scheme, and then to a different value of χ^2.

From Table 1, we know that the actual value of λ should be zero.[24] Therefore, when using the calibration scheme, we should recover a calibrated $\widehat{\lambda}$ close to zero, for most cases. Figure 4 shows four histograms of calibrated $\widehat{\lambda}$, one for each of the filtrations exposed in Table 2. Overall, the average volatility risk premium parameter is 0.671 when no options are used, and its standard error is 0.0969.[25] Therefore, the average parameter is statistically different from zero at a confidence level of 95%. When using one ATM option per day, the average calibrated parameter is -0.0282, and its standard error is 0.0051—still statistically different from zero at a confidence level of 95%. When using five and ten options, the average parameters are -0.0028 and -0.0017, respectively. Both are not statistically different from zero at a confidence level of 95% (standard error of 0.0018 in both cases).

All in all, using the approximated prices based only on past stock prices yields biased $\widehat{\lambda}$, which was expected as the option pricing formula used to obtain the calibrated parameter is biased. On the other hand, the use of option prices to obtain $\widehat{V}_t$ yield adequate, unbiased $\widehat{\lambda}$.

The results are fairly similar with the SVJ model. The average $\widehat{\lambda}$ is statistically different from zero when no options are used (at a 95% confidence level). When options are included in the filtering, the average $\widehat{\lambda}$ is closer to zero. One notable difference with the Heston case is that it requires more options (ten instead of five) to have an average $\widehat{\lambda}$ that is not statistically significant (still at a 95% confidence level).[26]

4 Empirical investigation

In this section, we perform an exercise similar to the one from the previous section, but now using real data. The goal is again to estimate the volatility risk premium parameter, λ, using the different estimates of V_t in Equations (4).[27]

[21] This approach has been used by Jacquier and Jarrow (1995), Madan et al. (1998) and Elliott et al. (1997), among others.

[22] Other loss functions were used and the results are robust to these other specifications.

[23] This calibration exercise is consistent with the method put forward by Broadie et al. (2007). In their paper, the authors find the risk-neutral parameters that minimize the absolute implied volatility errors, while keeping the physical parameters constant, similar to the methodology used in this section.

[24] In Aït-Sahalia and Kimmel (2007), they do not estimate this parameter because it is too cumbersome to estimate in their setting. The authors argue that there is a lot of uncertainty regarding λ in their setting, which is less of an issue in ours since we have a longer period and more option data.

[25] The standard error is computed here as the standard deviation of the sample of calibrated $\widehat{\lambda}$, divided by the square root of the sample size, which is 100.

[26] See Appendix A.5.2 and Figure 8 for more details.

[27] This time, in the filter, we use $\alpha = 0.20$. This value is consistent with Hurn et al.'s 2015 empirical results.

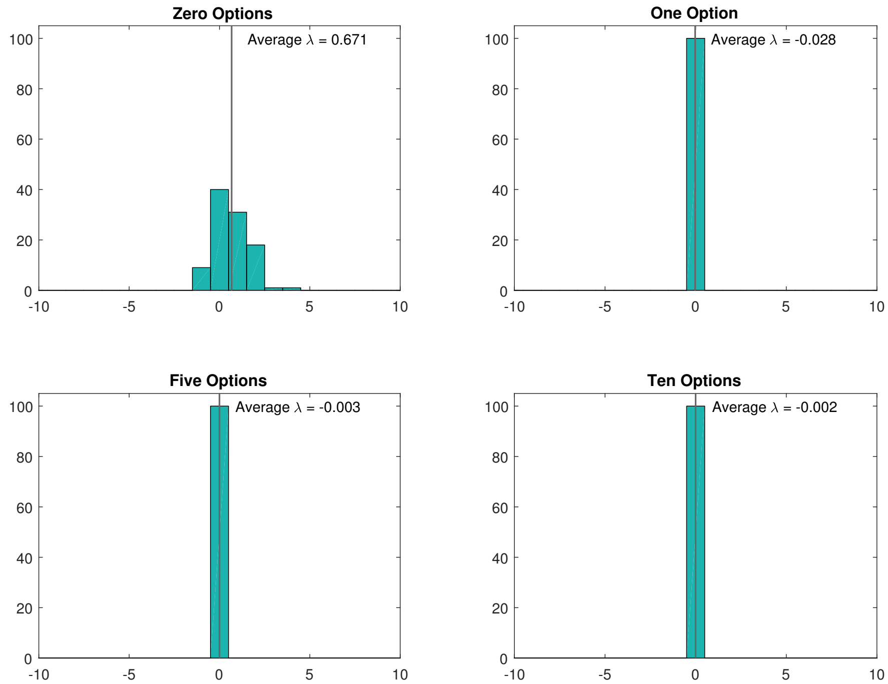

Figure 4: Histograms of calibrated variance premium parameter $\boldsymbol{\lambda}$ for the Heston model. This figure shows histograms of calibrated λ for four cases: zero options, one option, five options, and ten options. Using the asset prices and well-selected option prices (see Table 2), we apply the filtering methodology for each path, each day and each filtration to recover the filtered variance. We then compute the model with approximated prices, in order to calculate the χ^2 quantity of Equation (17). These two steps are done in turn, repeatedly, until we find the $\widehat{\lambda}$ that minimizes χ^2. This is achieved via numerical optimization by leaving parameter λ unrestrained.

Akin to the previous section, all the other parameters are set to those estimated by Aït-Sahalia and Kimmel (2007). The option data sample period employed to calibrate λ is from January 1996 to September 2003 in this section—consistent with the estimation sample used in Aït-Sahalia and Kimmel (2007).[28]

To obtain the best estimates $\widehat{V}_t$, we use daily log index returns obtained from the Center for Research in Security Prices (CRSP) and option prices obtained via OptionMetrics. The options selected in this study are European S&P 500 index option contracts (SPX). We use representative samples of options (similar to those introduced in Section 3.2) by restricting our attention to OTM or ATM options. Options with positive volume and bid price are included in the sample, as well as those satisfying the no-arbitrage conditions of Bakshi et al. (1997).[29]

Table 5 summarizes the various calibrated λ.[30] The estimated volatility risk premium $\widehat{\lambda}$ is -4.5661 when no options are used and -0.0724 when using only one ATM option. The last figure is rather close to the one hypothesized by Aït-Sahalia and Kimmel (2007); this is consistent with the fact that they use a single (synthetic) ATM short-dated option on each day in their estimation. The calibrated variance risk premium parameter is -0.5693, and -0.5671 when using five and ten options, respectively. Therefore, the bias illustrated in the previous section is also present in real data: the case using no options yields different estimators than those using five and ten options, a consequence of the propagation of the price bias on calibrated parameters. There is virtually no difference between the case with five and ten options, meaning that when the option price bias is reduced, λ is estimated more accurately.

[28]The option data of OptionMetrics starts in 1996. We therefore do not have the information for 1990–1995. Nonetheless, the sample we use is included in the one used by Aït-Sahalia and Kimmel (2007) and their parameters should be consistent with the data we employ in this study.

[29]We select OTM options within the OptionMetrics database that are the closest to (1) a moneyness level of 0.80, 0.85, 0.90, 0.95, 1.00, 1.05, 1.10, 1.15 and 1.20, and (2) a maturity of 30, 60, 90, 120 and 150 days. This yields a panel of 45 options on each day.

[30]Recall that in Aït-Sahalia and Kimmel (2007), this parameter is set to zero—it is not estimated.

Table 5: Calibrated variance risk premium parameter using the Heston model on the S&P 500 and data from 1996–2003.

	Zero Options	One Option	Five Option	Ten Options	ATM IV	VIX Index
$\widehat{\lambda}$	-4.5661	-0.0724	-0.5693	-0.5671	-2.9513	-0.3004

This table shows calibrated λ, using real data, for four cases: zero options, one option, five options and ten options. The sample period employed in this study is from January 1996 to September 2003 to calibrate λ. For the other parameters, we use those estimated by Aït-Sahalia and Kimmel (2007) using an MLE-based method, using only one "option" over a different sample period (i.e., 1990 to 2003); see Table 1 for more details. In addition to using the filtered variance obtained via the particle filter, we also use two other proxies for the instantaneous variance: the squared ATM implied volatility (IV) and the squared VIX index.

5 Other proxies

We now investigate the impact of using different proxies for the instantaneous variance in this section. Specifically, we repeat the experiments conducted in Sections 3 and 4 using ATM implied volatility and the VIX index. These two estimates are commonly used by both practitioners and scholars as proxies for the instantaneous variance by both practitioners and scholars.

Our first proxy for V_t is the squared short-term ATM Black and Scholes' (1973) implied volatility. This proxy is simply obtained by selecting from the panel of available options the one with a maturity closest to 30 days and a moneyness closest to 1, on a given day. The second proxy used in this section is the squared VIX index. To obtain the VIX index, we follow the methodology of the CBOE white paper (Chicago Board Options Exchange, 2009). This calculation is based on a panel of noisy simulated option prices (with a maturity of 30 days).

As in Section 3, we recover a surface of the option biases via a simulation exercise based on one hundred daily paths of the underlying asset over one year. For each path, we compute both our proxies and use them to obtain surfaces of option biases. Figure 5 shows the absolute bias as a function of the maturity of the option and its moneyness for the two proxies. While the option bias obtained by using the ATM implied volatility is mostly negative (except for short-term ATM options) and modest, the option bias is much worse and strictly positive when using the VIX index.[31] The surface obtained using the squared ATM implied volatility in lieu of V_t (left panel of Figure 5) is similar to that of the variance volga, and consistent with a large third term in Equation (10), which is a consequence of a low precision. On the other hand, when using the squared VIX, the bias surface is much more related to the vega, consistent with a large first-order estimate error term (i.e., the first term of Equation (10)), as shown in the right panel of Figure 5.[32]

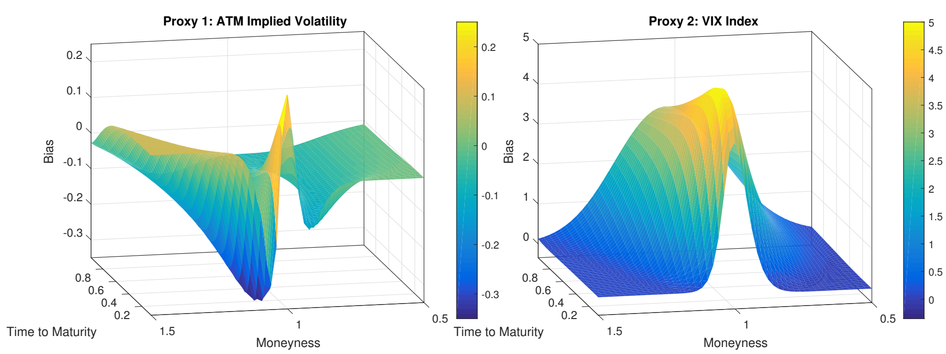

Figure 5: Option absolute bias for the two proxies, as a function of the time to maturity and the moneyness of the option for the Heston model. This figure shows the absolute option bias for both proxies. We compute the option price for various times to maturity and moneyness values using the approximation formula of Equation (4) based on the ATM implied volatility and the VIX index.

[31] Note that the scales on both surfaces are different as it would be impossible to see the leftmost surface if they both were on the same basis.

[32] For the affine model considered, the squared VIX is a linear function of V_t, e.g., $\text{VIX}^2 = a + bV_t$. Still, the factors a and b are functions of the model parameters (including λ). Computing these coefficients is not impossible in a two-stage estimation method (such as the one used in this study), although it makes things more complicated without really improving the inference. Indeed, as shown in Aït-Sahalia and Kimmel (2007), even if we account for a and b, we still obtain biased estimates of V_t.

The calibration exercise performed in Section 3 is also revisited: again using one hundred paths, we calibrate λ. However, this time, instead of using $\widehat{V}_t$ obtained via the particle filter, we use the two proxies: the squared ATM implied volatility and the squared VIX index. Figure 6 exhibits the histograms of calibrated $\widehat{\lambda}$ for the two proxies. For the ATM implied volatility, the average $\widehat{\lambda}$ is 0.3236, and this is statistically different from zero (standard error of 0.0793). The second proxy—the squared VIX—yields an average variance premium parameter of 6.1305, which is statistically different from zero.

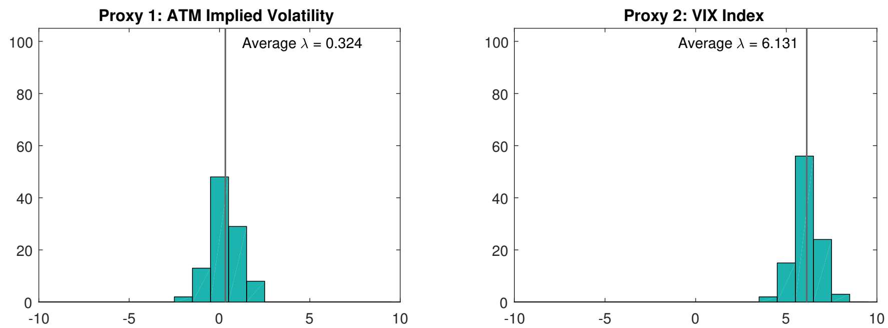

Figure 6: Histograms of calibrated variance premium parameter $\boldsymbol{\lambda}$ for the Heston model, continued. This figure shows histograms of calibrated λ for two cases: squared ATM implied volatility and squared VIX index. We compute the model with approximated prices, in order to calculate the χ^2 quantity of Equation (17). This step is repeated until we find the $\widehat{\lambda}$ that minimizes χ^2. This is achieved via numerical optimization, by leaving parameter λ unrestrained.

We have also redone the exercise of Section 4, but this time using the new proxies. The VIX index is extracted from the CBOE Indexes database. As shown in Table 5, the calibrated variance risk premium parameters are biased for both proxies: it is -2.9513 when using the squared ATM implied volatility and -0.3004 for the squared VIX index, i.e., not identical to the ones obtained with five options, $\mathcal{F}_t^5$, and ten options, $\mathcal{F}_t^{10}$.

6 Concluding remarks

In this paper, we have investigated the issue of using an estimate, also called the filtered value in the filtering context, to get option prices. We have shown that this approximation yields biased option prices only if the underlying values are used in the filter. If option values are also used in the filter, then the bias tends to disappear.

To gain more insight, we focused on the Heston model, which is one of the most elementary jump-diffusion models in the option pricing literature. We found that the bias is related to the first and second derivatives of the option pricing formula with respect to the instantaneous variance—the variance vega and volga. As expected, when using no options, the bias is positive when the variance volga is negative, and vice versa. On an absolute basis, the ATM options have the largest positive biases. ITM and OTM options have negative biases, especially those with short maturities. Even though absolute biases are not excessive for OTM options, a few pennies may have substantial consequences. For instance, an error of five cents might be considered large if the actual price of the option is about ten cents. In contrast, when using options in the filter, the bias is negative for all maturities and moneyness, but smaller (in absolute value): with one (ATM) option, the bias is still sizable, and when we include a larger number of options (e.g., five or ten), the bias virtually disappears. Similar results were obtained with the Bates (1996) SVJ model.

This bias could have drastic ramifications in calibration or estimation exercises that use option prices to recover model parameters, especially given that most of the recent studies use ATM and OTM options to infer the said parameters. Using simulated data, we find that the calibrated variance risk premium parameter is highly biased when no options are used in the filter; the bias is statistically significant from zero, at a confidence level of 95%. Nonetheless, the calibrated risk premium parameter is unbiased when using five or ten option prices. This is also the case when using real-world option prices: the volatility risk premium parameter is understated when no options are used. Indeed, using mistaken parameters obtained from the approximated pricing scheme can have many consequences: erroneous volatility surfaces and poor hedging strategies, for instance.

A Appendix

A.1 Option pricing formula

Commonly in jump-diffusions, a method based on partial differential equations is used to obtain option prices. With the latter, it is possible to price European vanilla options if the asset price dynamics are similar to Equations (1) and (2) under the risk-neutral measure. The key idea of this method is the calculation of the moment generating function of $\log(S_T)$, conditional on the latent factors.

Since our numerical illustration is based on the Heston model, we give, for convenience, the pricing equation for this specific jump-diffusion.

Proposition 2 *European Vanilla Option Price for the Heston Model. The price of a European call option with strike price K and maturity T is*

$$C_t(S_t, V_t) = S_t e^{-q(T-t)} P_1(S_t, V_t) - K e^{-r(T-t)} P_2(S_t, V_t), \tag{18}$$

where

$$P_1(S_t, V_t) = \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \mathrm{Re}\left(\frac{\exp\left(-\mathrm{i}u \log(K) - \log(S_t)\right) \varphi^{\mathbb{Q}}_{\log(S_T)|S_t,V_t}\left(ui+1, S_t, V_t\right)}{ui} \right) du,$$

$$P_2(S_t, V_t) = \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \mathrm{Re}\left(\frac{\exp\left(-\mathrm{i}u \log(K)\right) \varphi^{\mathbb{Q}}_{\log(S_T)|S_t,V_t}\left(ui, S_t, V_t\right)}{ui} \right) du.$$

and $\varphi^{\mathbb{Q}}_{\log(S_T)|S_t,V_t}(u, S_t, V_t) = \mathbb{E}^{\mathbb{Q}}[\exp\left(u \log(S_T)\right) \mid S_t, V_t]$ *is the moment generating function of* $\log(S_T)$ *conditional on* S_t *and* V_t*. More precisely,*

$$\varphi^{\mathbb{Q}}_{\log(S_T)|S_t,V_t}(u, S_t, V_t) = \exp\left(\mathcal{A}(u, T-t) + u\log(S_t) + \mathcal{B}(u, T-t) V_t\right),$$

where

$$\begin{aligned}
\mathcal{A}_V(u,\tau) &= (r-q)u\tau + \frac{\kappa^*\theta^*}{\sigma^2}\left((C_V - \gamma_V)\tau - 2\log\left(\frac{1 - D_V \exp\left(-\gamma_V \tau\right)}{1 - D_V}\right)\right), \\
\mathcal{B}_V(u,\tau) &= \frac{C_V - \gamma_V}{\sigma^2}\left(\frac{1 - \exp(-\gamma_V\tau)}{1 - D_V\exp(-\gamma_V\tau)}\right), \\
\gamma_V &= \sqrt{(\kappa^*)^2 + \left(u - u^2\right)\sigma^2}, \\
C_V &= \kappa^* - \rho\sigma u, \\
D_V &= \frac{C_V - \gamma_V}{C_V + \gamma_V}.
\end{aligned}$$

The approach relies on an inversion similar to that of Gil-Pelaez (1951). The price of a put option is obtained by the put-call parity relationship.

Proof. See Heston (1993) and Duffie et al. (2000). □

A.2 Proof of Proposition 1

Proof. Let $C \equiv C_t(S_t, V)$. The variance vega in the BS model is given by

$$\frac{\partial C}{\partial V} = \left(\frac{\partial C}{\partial\sqrt{V}}\right)\left(\frac{\partial\sqrt{V}}{\partial V}\right) = e^{-q(T-t)} S_t \phi(d_1) \frac{\sqrt{T-t}}{2\sqrt{V}},$$

where d_1 is given in Equation (12). For the previous equation, we know that

$$\frac{\partial^2 C}{\partial V^2} = \frac{\partial}{\partial V} e^{-q(T-t)} S_t \phi(d_1) \frac{\sqrt{T-t}}{2\sqrt{V}}$$

$$
\begin{aligned}
&= e^{-q(T-t)} S_t \sqrt{T-t} \left(\frac{1}{2\sqrt{V}} \frac{\partial}{\partial V} \phi(d_1) + \phi(d_1) \frac{\partial}{\partial V} \frac{1}{2\sqrt{V}} \right) \\
&= e^{-q(T-t)} S_t \sqrt{T-t} \left(\frac{1}{2\sqrt{V}} \frac{\partial \phi(d_1)}{\partial \sqrt{V}} \frac{\partial \sqrt{V}}{\partial V} + \phi(d_1) \frac{\partial}{\partial V} \frac{1}{2\sqrt{V}} \right) \\
&= e^{-q(T-t)} S_t \sqrt{T-t} \left(\frac{1}{4V} \phi(d_1) \left(\frac{d_1^2}{\sqrt{V}} - d_1 \sqrt{T-t} \right) - \phi(d_1) \frac{1}{4V^{3/2}} \right) \\
&= \frac{e^{-q(T-t)} S_t \sqrt{T-t} \phi(d_1)}{4V^{3/2}} \left(d_1^2 - d_1 \sqrt{V(T-t)} - 1 \right),
\end{aligned}
$$

since $\partial\phi(x)/\partial x = -x\phi(x)$. The volga for a put option is actually the same as the one for a call option, due to the put-call parity relationship. □

A.3 Variance vega and volga under the Heston model

In this section, we derive a semiclosed-form solution for the variance vega and volga under the Heston (1993) model. The variance vega is the first derivative of the option price with respect to the instantaneous variance and the variance volga is the second derivative of the option price with respect to the instantaneous variance.

Starting from Equation (18), the option pricing formula under the Heston model, we get

$$
\frac{\partial}{\partial V} C_t(S_t, V_t) = S_t e^{-q(T-t)} \left. \frac{\partial P_1(S_t, V)}{\partial V} \right|_{V=V_t} - K e^{-r(T-t)} \left. \frac{\partial P_2(S_t, V)}{\partial V} \right|_{V=V_t}
$$

for the variance vega, and

$$
\frac{\partial^2}{\partial V^2} C_t(S_t, V_t) = S_t e^{-q(T-t)} \left. \frac{\partial^2 P_1(S_t, V)}{\partial V^2} \right|_{V=V_t} - K e^{-r(T-t)} \left. \frac{\partial^2 P_2(S_t, V)}{\partial V^2} \right|_{V=V_t}
$$

for the variance volga.

The derivatives $\frac{\partial}{\partial V} P_1(S_t, V)$, $\frac{\partial}{\partial V} P_2(S_t, V)$, $\frac{\partial^2}{\partial V^2} P_1(S_t, V)$ and $\frac{\partial^2}{\partial V^2} P_2(S_t, V)$ can be calculated by computing the derivative of the integrand.[33] For instance, for the second derivative of $P_1(S_t, V)$ with respect to V, we have that

$$
\begin{aligned}
&\frac{\partial^2}{\partial V^2} P_1(S_t, V) \\
&= \frac{\partial^2}{\partial V^2} \left[\frac{1}{2} + \frac{1}{\pi} \int_0^\infty \mathrm{Re} \left(\frac{1}{u\mathrm{i}} \exp\left(-\mathrm{i}u \log(K) - \log(S_t)\right) \varphi^{\mathbb{Q}}_{\log(S_T)|S_t,V_t} \left(u\mathrm{i}+1, S_t, V\right) \right) du \right] \\
&= \frac{1}{\pi} \int_0^\infty \frac{\partial^2}{\partial V^2} \mathrm{Re} \left(\frac{1}{u\mathrm{i}} \exp\left(-\mathrm{i}u \log(K) - \log(S_t)\right) \varphi^{\mathbb{Q}}_{\log(S_T)|S_t,V_t} \left(u\mathrm{i}+1, S_t, V\right) \right) du \\
&= \frac{1}{\pi} \int_0^\infty \frac{\partial^2}{\partial V^2} \mathrm{Re} \Bigg(\frac{1}{u\mathrm{i}} \exp\left(-\mathrm{i}u \log(K) - \log(S_t)\right) \\
&\qquad\qquad \times \exp\left(\mathcal{A}\left(u\mathrm{i}+1, T-t\right) + \left(u\mathrm{i}+1\right) \log(S_t) + \mathcal{B}\left(u\mathrm{i}+1, T-t\right) V\right) \Bigg) du.
\end{aligned}
$$

Because $\exp(a + \mathrm{i}b) = \exp(a) \times (\cos b + \mathrm{i} \sin b)$, the last expression becomes

$$
\begin{aligned}
&\frac{1}{\pi} \int_0^\infty \frac{\partial^2}{\partial V^2} \mathrm{Re} \left(\frac{1}{u\mathrm{i}} \exp \left(\begin{array}{c} \mathrm{Re}\left[\mathcal{A}\left(u\mathrm{i}+1, T-t\right)\right] + \mathrm{Re}\left[\mathcal{B}\left(u\mathrm{i}+1, T-t\right)\right] V \\ +\mathrm{i}\mathrm{Im}\left[\mathcal{A}\left(u\mathrm{i}+1, T-t\right)\right] + \mathrm{i}\mathrm{Im}\left[\mathcal{B}\left(u\mathrm{i}+1, T-t\right)\right] V + \mathrm{i}u \log(S_t/K) \end{array} \right) \right) du \\
&= \frac{1}{\pi} \int_0^\infty \frac{\partial^2}{\partial V^2} \mathrm{Re} \left(\frac{1}{u\mathrm{i}} \left(\begin{array}{c} \exp\left(\mathrm{Re}\left[\mathcal{A}\left(u\mathrm{i}+1, T-t\right)\right] + \mathrm{Re}\left[\mathcal{B}\left(u\mathrm{i}+1, T-t\right)\right] V\right) \times \\ \left[\begin{array}{c} \cos\left(\mathrm{Im}\left[\mathcal{A}\left(u\mathrm{i}+1, T-t\right)\right] + \mathrm{Im}\left[\mathcal{B}\left(u\mathrm{i}+1, T-t\right)\right] V + u \log(S_t/K)\right) \\ +\mathrm{i} \sin\left(\mathrm{Im}\left[\mathcal{A}\left(u\mathrm{i}+1, T-t\right)\right] + \mathrm{Im}\left[\mathcal{B}\left(u\mathrm{i}+1, T-t\right)\right] V + u \log(S_t/K)\right) \end{array} \right] \end{array} \right) \right) du
\end{aligned}
$$

[33] Leibniz's integral rule is used here.

$$=\frac{1}{\pi}\int_0^\infty \frac{\partial^2}{\partial V^2}\mathrm{Re}\left(\frac{1}{u}\left(\begin{array}{c}\exp\left(\mathrm{Re}\left[\mathcal{A}\left(ui+1,T-t\right)\right]+\mathrm{Re}\left[\mathcal{B}\left(ui+1,T-t\right)\right]V\right)\times\\ \left[\begin{array}{c}-\mathrm{i}\cos\left(\mathrm{Im}\left[\mathcal{A}\left(ui+1,T-t\right)\right]+\mathrm{Im}\left[\mathcal{B}\left(ui+1,T-t\right)\right]V+u\log(S_t/K)\right)\\ +\sin\left(\mathrm{Im}\left[\mathcal{A}\left(ui+1,T-t\right)\right]+\mathrm{Im}\left[\mathcal{B}\left(ui+1,T-t\right)\right]V+u\log(S_t/K)\right)\end{array}\right]\end{array}\right)\right)du$$

$$=\frac{1}{\pi}\int_0^\infty \frac{\partial^2}{\partial V^2}\left(\frac{1}{u}\left(\begin{array}{c}\exp\left(\mathrm{Re}\left[\mathcal{A}\left(ui+1,T-t\right)\right]+\mathrm{Re}\left[\mathcal{B}\left(ui+1,T-t\right)\right]V\right)\\ \times\sin\left(\mathrm{Im}\left[\mathcal{A}\left(ui+1,T-t\right)\right]+\mathrm{Im}\left[\mathcal{B}\left(ui+1,T-t\right)\right]V+u\log(S_t/K)\right)\end{array}\right)\right)du.$$

Therefore,

$$\left.\frac{\partial^2}{\partial V^2}P_1(S_t,V)\right|_{V=V_t}$$

$$=\frac{1}{\pi}\int_0^\infty\left(\begin{array}{c}\frac{1}{u}\left(\begin{array}{c}\left(\mathrm{Re}\left[\mathcal{B}\left(ui+1,T-t\right)\right]^2-\mathrm{Im}\left[\mathcal{B}\left(ui+1,T-t\right)\right]^2\right)\\ \times\exp\left(\mathrm{Re}\left[\mathcal{A}\left(ui+1,T-t\right)\right]+\mathrm{Re}\left[\mathcal{B}\left(ui+1,T-t\right)V\right]\right)\\ \times\sin\left(\mathrm{Im}\left[\mathcal{A}\left(ui+1,T-t\right)\right]+\mathrm{Im}\left[\mathcal{B}\left(ui+1,T-t\right)V\right]+u\log(S_t/K)\right)\end{array}\right)\\ +\frac{1}{u}\left(\begin{array}{c}2\,\mathrm{Re}\left[\mathcal{B}\left(ui+1,T-t\right)\right]\mathrm{Im}\left[\mathcal{B}\left(ui+1,T-t\right)\right]\\ \times\exp\left(\mathrm{Re}\left[\mathcal{A}\left(ui+1,T-t\right)\right]+\mathrm{Re}\left[\mathcal{B}\left(ui+1,T-t\right)V\right]\right)\\ \times\cos\left(\mathrm{Im}\left[\mathcal{A}\left(ui+1,T-t\right)\right]+\mathrm{Im}\left[\mathcal{B}\left(ui+1,T-t\right)V\right]+u\log(S_t/K)\right)\end{array}\right)\end{array}\right)du.$$

The three other derivatives are obtained in a similar fashion. And again, the variance vega and volga for a call option are the same as the ones for a put option, thanks to the put-call parity relationship.

A.4 Particle filter

In this section, we briefly describe the sequential Monte Carlo method used to extract the density of the instantaneous variance and the expected (filtered) instantaneous variance. This implementation of the particle filter closely follows the work of Gordon et al. (1993). Specifically, we use the bootstrap filter. We also include the option prices in the filter by following the work of Hurn et al. (2015). The Error Specification 4 (ES4) of Hurn et al. (2015), a multiplicative pricing error specification centred at one, with a standard deviation of α, is used here; these authors argue that this specification is a "sensible choice because the mean value of the option price is not distorted by the choice of pricing error distribution."

This method is helpful as it allows us to recover both the density of V_t, $f_{\mathbb{P}}\left(V_t\,|\,\mathcal{F}_t\right)$, and the filtered value of the instantaneous variance, $\widehat{V}_t$.

A.5 Robustness test: Stochastic volatility with jumps instead of Heston

A.5.1 Option price bias

To test the robustness of our results, we add jumps to the stock price dynamics in the spirit of Bates' (1996) model. Specifically, we allow the stock price dynamics to experience jumps with a constant intensity λ_S. Using the notation of Equation (8), the SVJ model under the physical measure is given by the following stochastic differential equation:

$$dS_t=\left(r-q+\eta V_t+\xi-\lambda_S\mathbb{E}^{\mathbb{P}}\left[V_1\right]\right)S_{t^-}dt+\sqrt{V_t}S_{t^-}dW_{S,t}+S_{t^-}dJ_t,$$

where the process J_t is a compound Poisson process such that

$$J_t=\sum_{n=1}^{N_t}V_n,$$

N_t is a Poisson process with parameter λ_S, and $(1{+}V_n)$ follows a lognormal distribution such that $\log(1{+}V_n)\sim\mathcal{N}(\mu_S,\sigma_S^2)$. Following Hurn et al. (2015), we use the following dynamics under the risk-neutral measure:

$$dS_t=\left(r-q-\lambda_S^*\mathbb{E}^{\mathbb{Q}}\left[V_1\right]\right)S_{t^-}dt+\sqrt{V_t}S_{t^-}dW_{S,t}^{\mathbb{Q}}+S_{t^-}dJ_t^{\mathbb{Q}},$$

where

$$\xi=\left(\lambda_S-\lambda_S^*\right)\left(\exp\left(\mu_S+\sigma_S^2/2\right)-1\right).$$

Algorithm 1 Bootstrap filter.

1. For $i \in \{1, 2, ..., N\}$, the time t instantaneous variance $V_t^{(i)}$ is simulated from the proposal distribution

$$V_t^{(i)} = V_{t-1}^{(i)} + \kappa\left(\theta - V_{t-1}^{(i)}\right)h + \sigma\sqrt{V_{t-1}^{(i)}\,h}\,\varepsilon_{V,t}^{(i)},$$

where $\varepsilon_{V,t}^{(i)}$ is a standardized Gaussian random variable. This is equivalent to the well-known Euler-Maruyama discretization.

2. For $i \in \{1, 2, ..., N\}$, compute the option prices $\{C_{t,j}\left(S_t, V_t^{(i)}\right)\}_{j=1}^{M}$ based on S_t and $V_t^{(i)}$ from the semiclosed-form solution of Appendix A.1.

3. For $i \in \{1, 2, ..., N\}$, update the importance weights (up to a normalizing constant) to reflect how likely the simulated particles are with respect to the time t information:

$$\bar{\omega}_t^{(i)} = f\left(S_t \,\middle|\, S_{1:t-1}, V_{1:t}^{(i)}\right) \times f\left(\left\{C_{t,j}^{Market}\right\}_{j=1}^{M} \,\middle|\, S_{1:t}, V_{1:t}^{(i)}\right),$$

where $S_{1:t} = \{S_u : u \in \{1, 2, ..., t\}\}$. More precisely, using the linear interpolation for the drift and diffusion terms of Equations (8), the logarithm of the asset value satisfies

$$\begin{aligned}\log(S_t) &= \log(S_{t-1}) + (r-q)h + \left(\eta - \frac{1}{2}\right)\left(\frac{1}{2}\left(V_{t-1}^{(i)} + V_t^{(i)}\right)h\right) + \sqrt{\frac{1}{2}\left(V_{t-1}^{(i)} + V_t^{(i)}\right)h}\,\varepsilon_{S,t}^{(i)} \\ &= \log(S_{t-1}) + m + s\varepsilon_{S,t}^{(i)} \\ &= \log(S_{t-1}) + m + s\left(\rho\varepsilon_{V,t}^{(i)} + \sqrt{1-\rho^2}\,\varepsilon_{\perp,t}^{(i)}\right),\end{aligned}$$

where $\varepsilon_{S,t}^{(i)}$ and $\varepsilon_{\perp,t}^{(i)}$ are standardized Gaussian random variables, and $\varepsilon_{\perp,t}^{(i)}$ is independent of $\varepsilon_{V,t}^{(i)}$. Therefore, conditionally on the simulated instantaneous variances $V_{t-1}^{(i)}$ and $V_t^{(i)}$, and on the past asset value S_{t-1}, the time t asset value S_t is lognormally distributed,

$$f\left(S_t \,\middle|\, S_{t-1}, V_{t-1}^{(i)}, V_t^{(i)}\right) = \frac{1}{\sqrt{2\pi(1-\rho^2)}sS_t}\exp\left(-\frac{1}{2}\frac{\left(\log(S_t) - \log(S_{t-1}) - m - s\rho\varepsilon_{V,t}^{(i)}\right)^2}{(1-\rho^2)s^2}\right),$$

where $\varepsilon_{V,t}^{(i)}$ can be recovered from $V_{t-1}^{(i)}$ and $V_t^{(i)}$. Also, the option density is given by the following expression

$$f\left(\{C_{t,j}^{Market}\}_{j=1}^{M} \,\middle|\, S_{1:t}, V_{1:t}^{(i)}\right) = \prod_{j=1}^{M}\frac{1}{\sqrt{2\pi}\alpha C_{t,j}^{Market}}\exp\left(-\frac{1}{2}\frac{\left(\log\left(C_{t,j}^{Market}\right) - \log\left(C_{t,j}\left(S_t, V_t^{(i)}\right)\right) + \frac{1}{2}\alpha^2\right)^2}{\alpha^2}\right),$$

where α is a constant parameter related to the option prices' error.

4. For $i \in \{1, 2, ..., N\}$, compute the normalized weights

$$\omega_t^{(i)} = \frac{\bar{\omega}_t^{(i)}}{\sum_{k=1}^{N}\bar{\omega}_t^{(k)}}.$$

5. From normalized importance weights, you could compute the filtered variables

$$\widehat{V}_t = \sum_{i=1}^{N} V_t^{(i)}\omega_t^{(i)}$$

You could also keep the information about the weights and the simulated instantaneous variances to recover the (simulated) density of V_t, $f_{\mathbb{P}}(V_t \,|\, \mathcal{F}_t)$.

6. Resample the particles using the multinomial resampling.

   (a) Draw N particles from the current particle set from the empirical cumulative probability distribution and let $\left\{V_t^{(j_i)}\right\}_{i=1}^{N}$ denote the resulting underlying instantaneous variances once the resampling is performed.

   (b) Replace the current instantaneous variances with their resampled values:

$$V_t^{(i)} \leftarrow \quad V_t^{(j_i)}.$$

Using these $\mathbb{Q}-$dynamics, it is possible to find a closed-form solution for option prices,

$$C_t(S_t, V_t) = S_t e^{-q(T-t)}P_1'(S_t, V_t) - Ke^{-r(T-t)}P_2'(S_t, V_t),$$

where $P_z'(S_t, V_t)$ is similar to those obtained in Proposition 2 (see Duffie et al., 2000, for more details on the formula for the SVJ model).

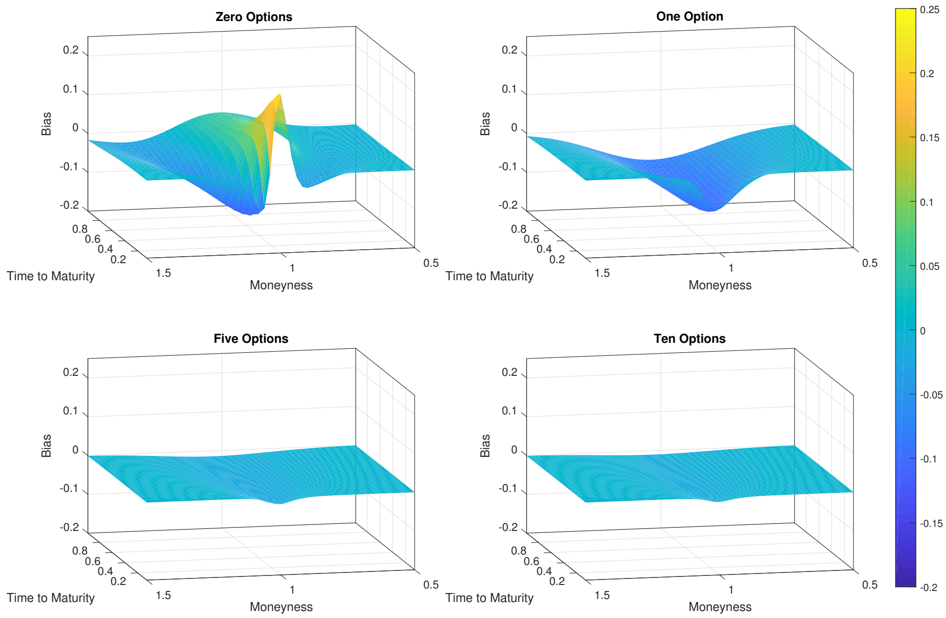

Figure 7: Option absolute bias as a function of the number of options used in the filter, the time to maturity and the moneyness of the option for the SVJ model. This figure shows the absolute bias for options, as a function of the number of options used in the filter (i.e., zero, one, five and ten). We compute the option price for various times to maturity and moneyness values using the approximation formula of Equations (4).

To assess the bias in the SVJ model, we use the parameters of Table 1, in combination with those found by Eraker et al. (2003) for the jump process: $\mu_S = -0.025862$, $\sigma_S = 0.040720$, $\lambda_S = 1.512$ and $\xi = 0$. Then, using these parameters, we repeat the exercise of Section 3.1: Figure 7 shows the absolute bias as a function of the maturity of the option and its moneyness for the four different instantaneous variance estimates (see Table 2). The results are highly similar to those of Figure 3: when no options are used, the bias is consistent with a large precision term in Equation (10). As soon as we start adding options, the bias becomes less significant—especially when the number of options used is larger, e.g., five or ten.[34]

A.5.2 Impact on calibrated parameter λ

Again, we know that the actual value of λ should be zero. Therefore, when using the calibration scheme, we should recover a calibrated $\widehat{\lambda}$ close to zero, for most cases. Figure 8 shows four histograms of calibrated $\widehat{\lambda}$ for the SVJ model, one for each of the filtration exposed in Table 2. The average volatility risk premium parameter is 0.4151 (standard error of 0.0811) when no options are used. When using one ATM option per day, the calibrated parameter is closer to zero, with an average of -0.0277 (standard error of 0.0040). When using five and ten options, the average parameters are -0.0075 and -0.0003, respectively (with standard errors of 0.0018 and 0.0016, respectively). Therefore, again, as we increase the number of options in the filtration, we get prices that are less biased, and therefore, a $\widehat{\lambda}$ parameter that is closer to zero. However, it takes more options than in the Heston case to recover the actual parameter, on average: with zero, one and five options, the average $\widehat{\lambda}$ is statistically different from zero, whereas the average is not statistically different from zero when we use ten options (at a confidence level of 95%).

At last, using the approximated prices based only on past stock prices yields biased $\widehat{\lambda}$, which was expected as the option pricing formula used to obtain the estimated parameter is biased. Contrarily, the use of option prices yield adequate, unbiased $\widehat{\lambda}$, especially when a large number of them is used.

[34] Additional results regarding the relative biases (similar to Tables 3 and 4) are given in Section SM.A.3 of the Supplementary Material.

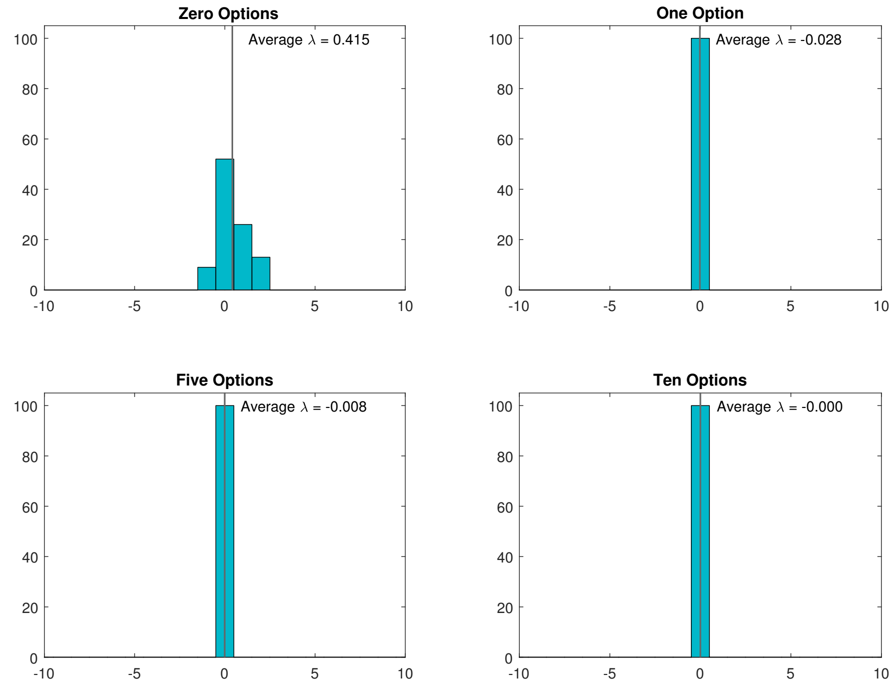

Figure 8: Histograms of calibrated variance premium parameter λ for the SVJ model. This figure shows histograms of calibrated λ for four cases: zero options, one option, five options, and ten options. Using the asset prices and well-selected option prices (see Table 2), we apply the filtering methodology for each path, each day and each filtration to recover the filtered variance. We then compute the model with approximated prices, in order to calculate the χ^2 quantity of Equation (17). These two steps are done in turn, repeatedly, until we find $\widehat{\lambda}$ that minimizes χ^2. This is achieved via numerical optimization, by leaving parameter λ unrestrained.

Supplementary material

SM.A Simulation experiments

SM.A.1 Impact of the maturity of the option

To assess the impact of the maturity of the option in the filtering procedure, we do over the simulation experiment of Section 3.1 using four different filtrations that each includes one option with a maturity of 15, 30, 90 and 150 days, respectively (see Table SM.1).

Table SM.1: The information structure used in the simulation experiment.

Maturity	Symbol	Description
15 Days	$\mathcal{F}_t^{1,\,[15]}$	This option has a maturity of 15 days and a moneyness of 1.00.
30 Days	$\mathcal{F}_t^{1,\,[30]}$	This option has a maturity of 30 days and a moneyness of 1.00.
90 Days	$\mathcal{F}_t^{1,\,[90]}$	This option has a maturity of 90 days and a moneyness of 1.00.
150 Days	$\mathcal{F}_t^{1,\,[150]}$	This option has a maturity of 150 days and a moneyness of 1.00.

All the filtrations include one option price on each day.

Figure SM.1 shows the option absolute bias as a function of the maturity of the option used in the filter, the time to maturity and the moneyness (in the spirit of Figure 3 of the manuscript). Interestingly, as the maturity of the ATM option used in the filter increases, the absolute error on ATM options become more negative. Using an option with maturity of 150 days in the filter, we obtain an absolute bias of about -0.25 for ATM options, which is about five times more than when we use an option with maturity of 15 or 30 days.

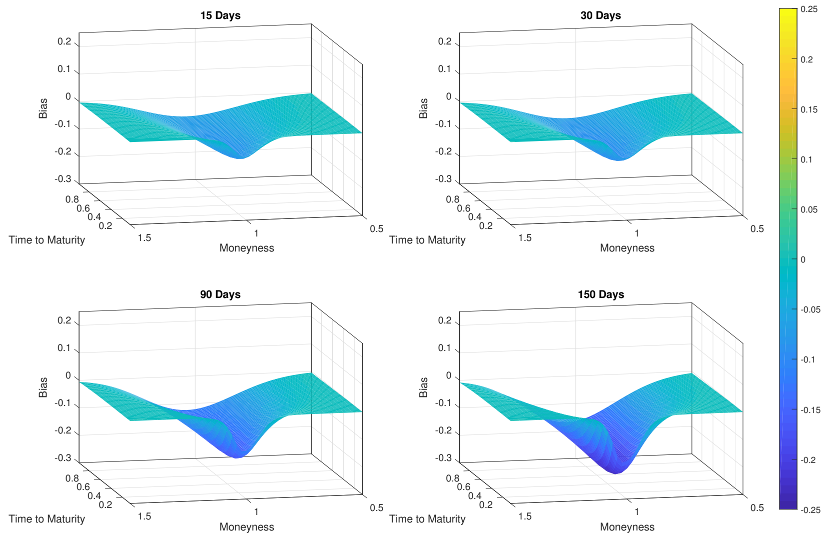

Figure SM.1: Option absolute bias as a function of the maturity of the option used in the filter, the time to maturity and the moneyness of the option for the Heston model. This figure shows the absolute bias for options, as a function of the maturity of the option used in the filter (i.e. 15 days, 30 days, 90 days, and 150 days).

SM.A.2 Impact on calibrated parameter λ

An interesting question relates to the impact of using different values of λ in our simulation study, i.e., different than zero. To assess the impact of this on our results of Section 3.1, we perform once more the simulation experiment using different values of λ. Figure SM.2 of this document—in the spirit of Figure 3—shows that changing the value of λ to 1 does not qualitatively change our results. Indeed, Figure 3 and Figure SM.2 are virtually identical. This stylized results stay the same for all the values of λ that we have investigated.

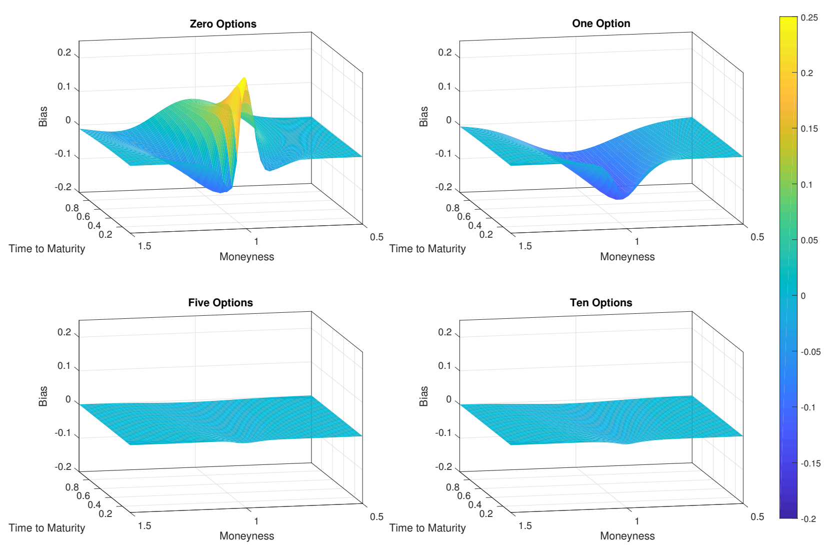

Figure SM.2: Option absolute bias as a function of the number of options used in the filter, the time to maturity and the moneyness of the option when λ is set to one for the Heston model. This figure shows the absolute bias for options, as a function of the number of options used in the filter (i.e. zero, one, five and ten).

SM.A.3 Option price bias for the SVJ model

In the case of the SVJ model, Tables SM.2 and SM.3 show the relative bias as a function of the maturity of the option and its moneyness for four different instantaneous variance estimates: using zero options, one option, five options, and ten options. These tables are in fact akin to Tables 3 and 4.

Table SM.2: Relative biases for calls as a function of time to maturity and moneyness for the SVJ model.

Panel A: Zero Options.					
Moneyness \ Maturity	$[0.0, 0.2]$	$[0.2, 0.4]$	$[0.4, 0.6]$	$[0.6, 0.8]$	$[0.8, 1.0]$
$[0.650, 0.750]$	0.000	0.005	0.012	0.017	0.020
$[0.750, 0.850]$	0.006	0.036	0.056	0.061	0.061
$[0.850, 0.950]$	0.116	0.211	0.208	0.184	0.162
$[0.950, 0.975]$	0.661	0.609	0.460	0.357	0.289
$[0.975, 1.000]$	1.295	0.892	0.610	0.452	0.355
$[1.000, 1.025]$	2.251	1.261	0.796	0.565	0.432
$[1.025, 1.050]$	3.232	1.728	1.025	0.700	0.522
$[1.050, 1.150]$	2.800	2.398	1.555	1.049	0.761
$[1.150, 1.250]$	-4.609	1.014	2.132	1.740	1.294
$[1.250, 1.350]$	-12.063	-4.945	0.304	1.751	1.711
Panel B: One Option.					
Moneyness \ Maturity	$[0.0, 0.2]$	$[0.2, 0.4]$	$[0.4, 0.6]$	$[0.6, 0.8]$	$[0.8, 1.0]$
$[0.650, 0.750]$	-0.000	-0.003	-0.005	-0.006	-0.006
$[0.750, 0.850]$	-0.004	-0.013	-0.017	-0.018	-0.018
$[0.850, 0.950]$	-0.037	-0.055	-0.054	-0.049	-0.044
$[0.950, 0.975]$	-0.146	-0.140	-0.113	-0.091	-0.076
$[0.975, 1.000]$	-0.255	-0.199	-0.148	-0.115	-0.093
$[1.000, 1.025]$	-0.439	-0.282	-0.193	-0.143	-0.113
$[1.025, 1.050]$	-0.746	-0.401	-0.253	-0.179	-0.137
$[1.050, 1.150]$	-1.508	-0.769	-0.436	-0.286	-0.207
$[1.150, 1.250]$	-3.979	-2.133	-1.087	-0.628	-0.411
$[1.250, 1.350]$	-4.797	-3.876	-2.236	-1.255	-0.768
Panel C: Five Options.					
Moneyness \ Maturity	$[0.0, 0.2]$	$[0.2, 0.4]$	$[0.4, 0.6]$	$[0.6, 0.8]$	$[0.8, 1.0]$
$[0.650, 0.750]$	-0.000	-0.000	-0.001	-0.001	-0.001
$[0.750, 0.850]$	-0.001	-0.002	-0.003	-0.003	-0.003
$[0.850, 0.950]$	-0.006	-0.010	-0.009	-0.009	-0.008
$[0.950, 0.975]$	-0.026	-0.025	-0.020	-0.016	-0.013
$[0.975, 1.000]$	-0.046	-0.035	-0.026	-0.020	-0.016
$[1.000, 1.025]$	-0.079	-0.050	-0.034	-0.025	-0.020
$[1.025, 1.050]$	-0.131	-0.071	-0.044	-0.031	-0.024
$[1.050, 1.150]$	-0.248	-0.131	-0.075	-0.050	-0.036
$[1.150, 1.250]$	-0.672	-0.346	-0.181	-0.107	-0.071
$[1.250, 1.350]$	-1.228	-0.708	-0.374	-0.210	-0.130
Panel D: Ten Options.					
Moneyness \ Maturity	$[0.0, 0.2]$	$[0.2, 0.4]$	$[0.4, 0.6]$	$[0.6, 0.8]$	$[0.8, 1.0]$
$[0.650, 0.750]$	-0.000	-0.000	-0.000	-0.001	-0.001
$[0.750, 0.850]$	-0.000	-0.001	-0.002	-0.002	-0.002
$[0.850, 0.950]$	-0.004	-0.007	-0.007	-0.006	-0.006
$[0.950, 0.975]$	-0.023	-0.021	-0.016	-0.013	-0.010
$[0.975, 1.000]$	-0.044	-0.031	-0.021	-0.016	-0.012
$[1.000, 1.025]$	-0.077	-0.044	-0.028	-0.020	-0.015
$[1.025, 1.050]$	-0.115	-0.060	-0.036	-0.025	-0.018
$[1.050, 1.150]$	-0.109	-0.087	-0.055	-0.037	-0.027
$[1.150, 1.250]$	-0.276	-0.027	-0.076	-0.062	-0.046
$[1.250, 1.350]$	-0.901	-0.324	-0.017	-0.057	-0.059

This table reports the average relative put option biases for each bin. The biases, calculated from Equation (16) for a large cross-section of options, are sorted into buckets of moneyness and maturity.

Table SM.3: Relative biases for puts as a function of time to maturity and moneyness for the SVJ model.

Panel A: Zero Options.					
Moneyness \ Maturity	$[0.0, 0.2]$	$[0.2, 0.4]$	$[0.4, 0.6]$	$[0.6, 0.8]$	$[0.8, 1.0]$
$[0.650, 0.750]$	0.374	2.180	1.893	1.428	1.091
$[0.750, 0.850]$	2.202	2.318	1.641	1.162	0.868
$[0.850, 0.950]$	2.913	1.932	1.223	0.846	0.631
$[0.950, 0.975]$	2.657	1.503	0.931	0.649	0.490
$[0.975, 1.000]$	2.236	1.263	0.799	0.567	0.434
$[1.000, 1.025]$	1.544	0.995	0.663	0.485	0.378
$[1.025, 1.050]$	0.821	0.734	0.535	0.408	0.326
$[1.050, 1.150]$	0.086	0.228	0.258	0.236	0.209
$[1.150, 1.250]$	-0.004	0.006	0.040	0.068	0.081
$[1.250, 1.350]$	-0.000	-0.002	0.001	0.010	0.021
Panel B: One Option.					
Moneyness \ Maturity	$[0.0, 0.2]$	$[0.2, 0.4]$	$[0.4, 0.6]$	$[0.6, 0.8]$	$[0.8, 1.0]$
$[0.650, 0.750]$	-2.337	-1.221	-0.723	-0.483	-0.350
$[0.750, 0.850]$	-1.614	-0.837	-0.503	-0.342	-0.252
$[0.850, 0.950]$	-0.921	-0.506	-0.320	-0.225	-0.170
$[0.950, 0.975]$	-0.588	-0.345	-0.228	-0.166	-0.129
$[0.975, 1.000]$	-0.441	-0.282	-0.194	-0.144	-0.113
$[1.000, 1.025]$	-0.301	-0.222	-0.161	-0.123	-0.099
$[1.025, 1.050]$	-0.190	-0.171	-0.132	-0.104	-0.085
$[1.050, 1.150]$	-0.046	-0.073	-0.072	-0.065	-0.057
$[1.150, 1.250]$	-0.003	-0.012	-0.021	-0.025	-0.026
$[1.250, 1.350]$	-0.000	-0.001	-0.004	-0.007	-0.009
Panel C: Five Options.					
Moneyness \ Maturity	$[0.0, 0.2]$	$[0.2, 0.4]$	$[0.4, 0.6]$	$[0.6, 0.8]$	$[0.8, 1.0]$
$[0.650, 0.750]$	-0.398	-0.206	-0.123	-0.083	-0.060
$[0.750, 0.850]$	-0.269	-0.142	-0.087	-0.059	-0.044
$[0.850, 0.950]$	-0.158	-0.088	-0.056	-0.039	-0.030
$[0.950, 0.975]$	-0.104	-0.061	-0.040	-0.029	-0.022
$[0.975, 1.000]$	-0.079	-0.050	-0.034	-0.025	-0.020
$[1.000, 1.025]$	-0.054	-0.039	-0.028	-0.021	-0.017
$[1.025, 1.050]$	-0.033	-0.030	-0.023	-0.018	-0.015
$[1.050, 1.150]$	-0.008	-0.012	-0.012	-0.011	-0.010
$[1.150, 1.250]$	-0.001	-0.002	-0.003	-0.004	-0.004
$[1.250, 1.350]$	-0.000	-0.000	-0.001	-0.001	-0.002
Panel D: Ten Options.					
Moneyness \ Maturity	$[0.0, 0.2]$	$[0.2, 0.4]$	$[0.4, 0.6]$	$[0.6, 0.8]$	$[0.8, 1.0]$
$[0.650, 0.750]$	-0.019	-0.073	-0.066	-0.050	-0.038
$[0.750, 0.850]$	-0.074	-0.082	-0.058	-0.041	-0.030
$[0.850, 0.950]$	-0.104	-0.068	-0.043	-0.030	-0.022
$[0.950, 0.975]$	-0.092	-0.052	-0.033	-0.023	-0.017
$[0.975, 1.000]$	-0.076	-0.044	-0.028	-0.020	-0.015
$[1.000, 1.025]$	-0.053	-0.035	-0.023	-0.017	-0.013
$[1.025, 1.050]$	-0.029	-0.026	-0.019	-0.014	-0.011
$[1.050, 1.150]$	-0.003	-0.008	-0.009	-0.008	-0.007
$[1.150, 1.250]$	-0.000	-0.000	-0.001	-0.002	-0.003
$[1.250, 1.350]$	-0.000	-0.000	-0.000	-0.000	-0.001

This table reports the average relative put option biases for each bin. The biases, calculated from Equation (16) for a large cross-section of options, are sorted into buckets of moneyness and maturity.

References

Aït-Sahalia, Y., and R. Kimmel. 2007. Maximum likelihood estimation of stochastic volatility models. Journal of Financial Economics 83:413–452.

Bakshi, G., C. Cao, and Z. Chen. 1997. Empirical performance of alternative option pricing models. Journal of Finance pp. 2003–2049.

Bardgett, C., E. Gourier, and M. Leippold. 2016. Inferring volatility dynamics and risk premia from the S&P 500 and VIX markets. Working paper .

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.

Black, F., and M. Scholes. 1973. The pricing of options and corporate liabilities. Journal of Political Economy 81:637–654.

Broadie, M., M. Chernov, and M. Johannes. 2007. Model specification and risk premia: evidence from futures options. Journal of Finance 62:1453–1490.

Carr, P., and L. Wu. 2010. Stock options and credit default swaps: A joint framework for valuation and estimation. Journal of Financial Econometrics 8:409–449.

Chicago Board Options Exchange. 2009. The CBOE volatility index—VIX. White Paper .

Christoffersen, P., S. Heston, and K. Jacobs. 2009. The shape and term structure of the index option smirk: Why multifactor stochastic volatility models work so well. Management Science 55:1914–1932.

Christoffersen, P., K. Jacobs, and K. Mimouni. 2010. Volatility dynamics for the S&P500: 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.

Duffie, D., J. Pan, and K. Singleton. 2000. Transform analysis and asset pricing for affine jump-diffusions. Econometrica 68:1343–1376.

Elliott, R. J., C. H. Lahaie, and D. B. Madan. 1997. Filtering derivative security valuations from market prices. In M. A. Dempster and S. R. Pliska (eds.), Mathematics of Derivative Securities, chap. 10, pp. 141–162. Cambridge University Press.

Eraker, B. 2001. MCMC analysis of diffusion models with application to finance. Journal of Business & Economic Statistics 19:177–191.

Eraker, B. 2004. Do stock prices and volatility jump? Reconciling evidence from spot and option prices. Journal of Finance 59:1367–1404.

Eraker, B., M. Johannes, and N. Polson. 2003. The impact of jumps in volatility and returns. Journal of Finance 58:1269–1300.

Gil-Pelaez, J. 1951. Note on the inversion theorem. Biometrika 38:481–482.

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.

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.

Huang, J.-Z., and L. Wu. 2004. Specification analysis of option pricing models based on time-changed Lévy processes. Journal of Finance 59:1405–1439.

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., and R. Jarrow. 1995. Dynamic evaluation of contingent claim models: An analysis of model error. 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.

Jones, C. S. 1998. Bayesian estimation of continuous-time finance models. Working paper .

Jones, C. S. 2003. The dynamics of stochastic volatility: Evidence from underlying and options markets. Journal of Econometrics 116:181–224.

Kalman, R. E. 1960. A new approach to linear filtering and prediction problems. Journal of Basic Engineering 82:35–45.

Madan, D. B., P. P. Carr, and E. C. Chang. 1998. The variance gamma process and option pricing. European Finance Review 2:79–105.

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.

Rebonato, R. 2005. Volatility and correlation: The perfect hedger and the fox. John Wiley & Sons.

Santa-Clara, P., and S. Yan. 2010. Crashes, volatility, and the equity premium: lessons from S&P 500 options. Review of Economics and Statistics 92:435–451.

Sepp, A. 2008. Pricing options on realized variance in Heston model with jumps in returns and volatility. Journal of Computational Finance 11:33–70.