

American Option Valuation with Particle Filters*

Bhojnarine R. Rambharat

Abstract A method to price American-style option contracts in a limited information framework is introduced. The pricing methodology is based on sequential Monte Carlo techniques, as presented in Doucet, de Freitas, and Gordon's text *Sequential Monte Carlo Methods in Practice*, and the least-squares Monte Carlo approach of Longstaff and Schwartz (Rev Financ Stud 14:113–147, 2001). We apply this methodology using a risk-neutralized version of the square-root mean-reverting model, as used for European option valuation by Heston (Rev Financ Stud 6:327–343, 1993). We assume that volatility is a latent stochastic process, and we capture information about it using particle filter based “summary vectors.” These summaries are used in the exercise/hold decision at each time step in the option contract period. We also benchmark our pricing approximation against the full-state (observable volatility) result. Moreover, posterior inference, utilizing market-observed American put option prices on the NYSE Arca Oil Index, is made on the volatility risk premium, which we assume is a constant parameter. Comparisons on the volatility risk premium are also made in terms of time and observability effects, and statistically significant differences are reported.

Keywords American options • Latent • Monte Carlo • Optimal stopping • Optimization • Particle filter • Posterior inference • Risk premium • Volatility

MSC code: 62F15, 62L15, 91G20, 91G60, 91G70

**Disclaimer:* The views expressed in this paper are solely those of the author and neither reflect the opinions of the OCC nor the U.S. Department of the Treasury.

B.R. Rambharat (✉)

Department of Treasury, Office of the Comptroller of the Currency (OCC), 250 E Street SW, Washington, DC 20219, USA

e-mail: ricky.rambharat@occ.treas.gov

1 Introduction

The valuation of American-style options is an evergreen area of research in quantitative finance. The solution to the American option valuation problem has widespread, practical applications, especially as regards the products offered on major financial exchanges throughout the world. The distinguishing feature of an American-style financial contract is the fact that it can be exercised at any point from its inception date to its expiration date. This is in stark contrast to a European-style option, which can only be exercised at its expiration. Consequently, an American option poses a significantly more difficult valuation problem relative to its European counterpart.

The value of an American-style (early-exercise) option, guided by the “fundamental theorem of no-arbitrage pricing” due to [31], is found by calculating the discounted expectation of the relevant payoff function under a risk-neutral measure, assuming that exercise/hold decisions are made to maximize the payoff function. The option, or derivative, is typically based on underlying price series (e.g., equity price, interest rate, index value, etc.) whose random fluctuations are commonly modeled using stochastic processes. The main difficulty with American option pricing is obtaining a reliable estimate of the hold (continuation) value. Assuming that all random factors affecting the underlying price series are fully observable, the price of an American option is computed by solving an optimal stopping problem using the principles of dynamic programming as set forth in [3]. Although observability of all sources of randomness is an unrealistic assumption in real-world financial markets, a majority of all research on American-style option valuation starts from the assumption that all sources of randomness are fully observable, and hence, can be encompassed in the state-space of the associated optimal stopping problem. A few key, although non-exhaustive, computational references on American option pricing include [5, 7, 10, 21, 28, 68]. The Monte Carlo based valuation algorithm of [46] has found significant practical application since its introduction, and we use it extensively in what follows. A solid theoretical treatment on American option valuation is given in [51].

The notion of observability finds a ripe application with stochastic volatility models since it is arguably the case that volatility is latent. The celebrated work of [4] provides a solution to the European option pricing problem in an arbitrage-free, *constant volatility* framework. Indeed, this seminal paper also provides a way to estimate the volatility based on observed option prices, however, this assumes that the volatility parameter remains static throughout the life of the option contract. Notwithstanding, empirical findings, such as the volatility “smile” or “smirk,” suggest that volatility is actually not constant. Since the work of [4], a literature on stochastic volatility option pricing emerged and some key examples include, among others, [26, 33, 35, 67]. Since European options have a fixed exercise date, which is known at the inception of the option, volatility can be effectively “averaged out” if it is indeed treated as a stochastic process, and this is confirmed in the aforementioned references.

The case of American options under stochastic volatility presents additional challenges, both theoretically and computationally. If it is assumed that volatility is stochastic and observable, then as noted above, the price of an American option can be computed utilizing the standard principles of dynamic programming. A few examples of work treating the valuation of American options in a stochastic volatility setting include [6, 15, 25, 26, 30, 48, 69, 71, 76]. A majority of these references include the share price and volatility as state variables in the pricing algorithm, thus assuming that the volatility process is observable. One exception is [26] where the authors provide a useful (and practical) correction to the constant volatility option price using the implied volatility surface within a “fast mean–reversion” framework. In [57], an approximate grid–based solution is proposed for the “limited information” optimal stopping problem using an illustrative stochastic volatility model, where volatility follows a latent, geometric mean–reverting process (i.e., a “Schwartz Type I process” due to [65]), or equivalently, where the natural logarithm of volatility follows an Ornstein–Uhlenbeck (OU) process. Additionally, the work in [58] proposes a Monte Carlo based pricing methodology that combines the least–squares Monte Carlo (LSM) algorithm of [46] along with a sequential Monte Carlo filter as presented in [23] in order to find the optimal pricing solution using the above noted log–stochastic volatility model.

Additional recent work by researchers have studied the problem of latent processes in an optimal stopping or portfolio optimization framework. The work in [19] addresses partially observed stochastic volatility in a portfolio optimization problem, while the analysis in [54] illustrates a quantization method to solve optimal stopping problems under a partial information framework. Moreover, [47] proposes theoretical and numerical results using a similar LSM/particle filter approach for optimal stopping problems with partial information, and applied to a more generic diffusion modeling specification.

The present work adds to the research concerning the problem of American option valuation in the presence of *unobserved* sources of randomness, specifically volatility, that affect the underlying price series on which the option contract is written. Posterior inference on these unobservables is accomplished using sequential Monte Carlo (or particle filtering) methodology. Our goal is to find the optimal exercise rule for American options where the underlying price series is governed by a stochastic volatility model, assuming that volatility is unobservable. Particle filtering methods have been studied in the context of option pricing in [37] and [47]. Moreover, the neural network analysis in [27] can potentially be adapted to American option pricing as was done in [70].

The current analysis uses the regression–based LSM algorithm, along with particle filtering techniques, to solve an optimal stopping problem with partial information, and also implement a rigorous empirical analysis in order to apply the proposed method to market–observed American option prices. Moreover, we aim to make posterior inference on the market price of volatility risk in light of the effects of “observed versus unobserved stochastic volatility” as well as the effects of “time to expiration.” We extend the earlier research in [58] and [57] along a few key themes. First, the actual stochastic volatility modeling framework is based on

the square-root mean-reverting model as analyzed in [33], which directly imposes positivity constraints on the actual volatility process. This model appears in the seminal work of [16] in the context of modeling interest rates, and it is commonly used for stochastic volatility modeling in practical settings. We also benchmark our approximations to the optimal American option price obtained from the full observation case.

This paper is organized according to the following sections. Section 2 discusses the valuation problem in a stochastic volatility framework with emphasis on the latency of the volatility process. We also discuss the specific stochastic volatility model that we use, which is based on the work of [33]. Next, Sect. 3 describes the steps of our particle filtering based pricing algorithm. Subsequently, in Sect. 4, we assess the quality of our American option price approximations using a numerical benchmark analysis. Section 5 presents the results of an empirical application of our pricing methodology to market-observed American-style put options on the NYSE Arca Oil Index. Finally, Sect. 6 offers concluding remarks and discusses outstanding issues for future research.

2 Valuation Framework

Stochastic volatility models are arguably one of the most well-studied areas within the research on option valuation. A stochastic volatility model adds more flexibility relative to other modeling frameworks for the purpose of option pricing. The evidence in the extant literature shows that stochastic volatility models are able to offer insights into “empirical peculiarities,” such as smile/smirk effects, that constant volatility models are not able to capture. The research in [35] proposes a stochastic volatility model based on geometric Brownian motion, however, later studies favored the use of mean-reverting models, as is evidenced by the works of [67] and [33].

Several challenges arise with stochastic volatility models, especially as regards option valuation. One of the main challenges is the choice of a risk-neutral pricing measure. An additional issue concerns the accuracy of the simulation methodology used for option pricing. Concerning simulation, there is typically a trade-off between pricing accuracy and computational expense. The observability of volatility is yet another challenge associated with stochastic volatility models, and we show in this work how this has significant implications for American options. The aforementioned issues are not meant to be exhaustive, however, they are critical to stochastic volatility models and we elaborate on each in turn.

2.1 A Risk-Neutral Stochastic Volatility Model

A risk-neutral version of the square-root mean-reverting stochastic volatility model is

$$dS_t = rS_t dt + \sqrt{V_t}S_t \left(\sqrt{1 - \rho^2} dW_1(t) + \rho dW_2(t) \right), \quad (1)$$

$$dV_t = (\alpha + \lambda\gamma) \left(\frac{\alpha\beta}{\alpha + \lambda\gamma} - V_t \right) dt + \gamma \sqrt{V_t} dW_2(t), \quad (2)$$

where S_t denotes the underlying observed process (e.g., share price, index value, etc.), r is the risk-free rate of return, V_t is the square of the volatility (or stochastic variance of the process), $\alpha + \lambda\gamma$ is the rate of mean-reversion of V_t , $\frac{\alpha\beta}{\alpha + \lambda\gamma}$ is the level of mean-reversion of V_t , γ is the volatility of V_t , ρ is the correlation between the two Brownian motions, $\sqrt{1 - \rho^2} dW_1(t) + \rho dW_2(t)$ and $dW_2(t)$, and finally λ is a parameter that quantifies the volatility risk premium or the “market price of volatility risk.” (If we set $\alpha^* = \alpha + \lambda\gamma$ and $\beta^* = \frac{\alpha\beta}{\alpha^*}$, then the expression for dV_t can be written in the usual form that expresses a square-root mean-reverting process.)

The formal price of an American-style option, with respect to a risk-neutral measure parameterized by λ , is given by

$$\pi_t = \sup_{\tau \in \mathcal{T}} E_\lambda [e^{-r\tau} g(S_\tau, \tau) | \mathcal{D}_t], \quad (3)$$

where t is the current time, \mathcal{D}_t is the filtration generated by the available *observed* data, τ is a random stopping time at which an exercise decision is made, and \mathcal{T} is the set of all possible stopping times with respect to the filtration \mathcal{D}_t . Additionally, the function $g(S_\tau, \tau)$ is the payoff of the option. For instance, in the case of a call option, $g(S_\tau, \tau) = \max\{S_\tau - K, 0\}$, and in the case of a put option, $g(S_\tau, \tau) = \max\{K - S_\tau, 0\}$, where in both cases K is known as the strike price.

Regarding volatility, since it is not a traded asset, it is not possible to perfectly hedge away all of the risks associated with it. Therefore, valuation takes place in an “incomplete market” framework. The quantity λ is not uniquely determined in the valuation problem. There are potentially numerous ways to specify λ . For the purpose of our analysis, we regard λ as a constant and we later demonstrate how to estimate it using market-observed option prices.

In (1) and (2), the rate and level of mean-reversion are parameterized under the risk-neutralized pricing measure. If we set $\lambda = 0$, we will find that α and β characterize the rate and level of mean-reversion, respectively, which is indeed the case under the statistical or “real-world” measure. Under the statistical measure, the parameters of the V_t process satisfy constraints that ensure positivity of the process, namely, $\alpha\beta > 0$ and $\gamma^2 < 2\alpha\beta$ (see [16] for details).¹ Moreover, as λ becomes more negative, the drift in the volatility increases. The increase in the volatility drift increases spot volatility, which generally increases option prices.

A full theoretical treatment of pricing in incomplete markets, or under stochastic volatility models in general, is beyond the scope of this work. The research by [72]

¹Note that these positivity constraints for the square-root mean-reverting model are also satisfied under the risk-neutral measure.

and [20] present volatility as an underlying asset on which option contracts are written. The work in [41] presents some very important results on pricing derivatives in incomplete markets. Additional useful references on the topic of pricing in incomplete markets include, but are not limited to, [9, 32, 34, 49]. In this analysis, we will regard volatility as an asset that is not traded, and moreover, we regard it as a latent stochastic process. We now turn to a discussion of how to simulate the process in (1) and (2).

2.2 Simulation Methodology

Accurate and efficient simulation from a given model is pivotal to numerical analysis of historical price series or market observed option prices. This point is especially crucial for stochastic volatility models since we often rely on approximate simulation methodology to understand them. In [29], a very thorough discussion is given of Monte Carlo simulation methods in quantitative finance, including applications to American options. The pricing methodology that we use in this paper is based on the simulation approach of [46] as well as the earlier work of [11]. Another Monte Carlo based pricing approach that could be used for American options is the stochastic mesh method that is also discussed in [29].

The exact solution for the price process in (1) and (2) is

$$S_t = S_{t-\Delta} \exp \left\{ r\Delta - \frac{1}{2} \int_{t-\Delta}^t V_s ds + \sqrt{1-\rho^2} \int_{t-\Delta}^t \sqrt{V_s} dW_1(s) + \rho \int_{t-\Delta}^t \sqrt{V_s} dW_2(s) \right\}, \quad (4)$$

where $\sqrt{V_s}$ is the volatility at time s , Δ is the time step between time $t - \Delta$ and t , and all other parameters are as previously defined. The exact solution for the V_t process, conditional on $V_{t-\Delta}$ is

$$V_t = V_{t-\Delta} + \alpha\beta\Delta - (\alpha + \lambda\gamma) \int_{t-\Delta}^t V_s ds + \gamma \int_{t-\Delta}^t \sqrt{V_s} dW_2(s). \quad (5)$$

As discussed in [16], V_t given V_s , $s < t$, is distributed as a scaled Non-central Chi-squared random variable.

Definition 2.1. Let Z_i be independent and identically distributed Normal random variables, where $E(Z_i) = \mu_i$ and $Var(Z_i) = 1$, for $i = 1, \dots, k$. Let $\lambda = \sum_{i=1}^k \mu_i^2$, and set $\chi = \sum_{i=1}^k Z_i^2$. Then χ is distributed as a Non-central Chi-squared random variable with k degrees of freedom and non-centrality parameter λ .

The above definition is motivated from [38]. Additional details of the square-root mean-reverting model are also explained in [8], where it is stated that the Non-central Chi-squared is essentially a central Chi-squared with random degrees of freedom. Therefore, a simple way to simulate a draw, χ , from a Non-central Chi-squared random variable with k degrees of freedom and non-centrality parameter λ would be to use the steps in the following routine.

Routine 1: Non-central Chi-Squared simulation

1. Generate $\xi \sim \text{Poisson}$ with mean $\lambda/2$.
2. Generate $\chi \sim \text{central Chi-squared}$ with degrees of freedom (d.f.) equal to $k + 2\xi$.
3. Return χ .

The exact mapping of the parameter set of the V_t process in (2) (i.e., mean-reversion rate, mean-reversion level, and volatility-of-volatility) to the parameters of the aforementioned Poisson and Chi-squared random variables is intricate, however, the details are available in [16] or [38]. In [8], an exact solution to (4) is provided using Fourier methods, along with numerical integration/optimization techniques. The full details of our core simulation procedure for (S_t, V_t) is outlined in the next routine.

Routine 2: Price simulation

1. Initialize: let $\Delta = \text{time step}$, and let (S_0, V_0) be initial values at $t = 0$.
2. Parameter definitions (see (1) – (2)).
 - Set $\alpha^* = \alpha + \lambda\gamma$ and $\beta^* = \frac{\alpha\beta}{\alpha^*}$.
 - Set $a = \frac{2\alpha^*}{\gamma^2(1-e^{-\alpha^*\Delta})}$.
 - Set $b = a \cdot e^{-\alpha^*\Delta} \cdot V_{t-\Delta}$.
 - Set $c = \frac{2\alpha^*\beta^*}{\gamma^2} - 1$.
3. V_t simulation (see Routine 1 and (5)).
 - Generate $\xi \sim \text{Poisson}$ with mean u .
 - Generate $\chi \sim \text{Chi-squared}$ with d.f. equal $2c + 2 + 2\xi$.
 - Set $V_t = \frac{\chi}{2a}$.

4. S_t simulation (see (4)).

- Solve for $\int_{t-\Delta}^t \sqrt{V_s} dW_2(s)$ in (5) using V_t simulation output.
- Generate $\int_{t-\Delta}^t \sqrt{V_s} dW_1(s) \sim \text{Normal}$ with mean 0 and variance $E \left[\int_{t-\Delta}^t V_s ds \right]$ in (4) using properties of the stochastic integral and, specifically, the Itô isometry.
- Generate S_t according to (4).

While we simulate draws from V_t conditional on $V_{t-\Delta}$ exactly, we use a first-order approximation to evaluate $\int_{t-\Delta}^t V_s ds$ when we generate S_t conditional on $S_{t-\Delta}$. Indeed, all of our approximations are confined to this Euler approximation, which we use in order to minimize computational cost. The above conclusions regarding the stochastic integral are based on results found in [40].

2.3 Latent Volatility

As noted above, stochastic volatility models create an incomplete market valuation framework since volatility is not a traded asset. Specification of the volatility risk premium parameter, λ , is required for the purpose of option valuation. An additional challenge associated with stochastic volatility models involves observability. It is unrealistic to assume that all information in the market is observable. Regarding volatility, one could build a pricing model based on the assumption of full observability, and then estimate the volatility *implied* from market observed option prices. For instance, [18] estimates the volatility smile from American options. Moreover, [60] filter the spot volatility in a stochastic volatility model from observed option prices.

The research treating stochastic volatility as a latent quantity, upon which posterior inference is made, is limited. The work in [73] provides a theoretical framework for combining sequential decisions with posterior inference under certain conditions. The problem of combining latent stochastic volatility with American option pricing creates additional layers of complexity. Note that the notion of observed versus latent volatility is *not* an issue for European-style options. Since the exercise date is fixed, we can effectively integrate over the average distribution of the volatility from the inception time t to the expiration time T . This is essentially a conditional Monte Carlo approach, and an example of how it can be implemented in a European option pricing framework is discussed in [35].

The limited information aspect of latent stochastic volatility results in more challenges for American-style options due to the inherent sequential nature of the pricing problem. If volatility were observed, the owner of an American option would

use information on both S_t and V_t in order to make an exercise/hold decision. The optimal exercise boundary would also be a function of both the price and volatility. Assuming that volatility is latent, the owner of an American option only has information on the observed process S_t and must make posterior inference on V_t conditional on the observed process. The posterior filtering distribution, which we denote by $p(V_t|s_1, \dots, s_t)$, where s_1, \dots, s_t are realized values of the observed process, increases in dimension, as a measure of probability, as more data are observed. Consequently, the dynamic programming pricing algorithm will be effectively based on an infinite-dimensional state-space.

If the distribution $p(V_t|s_1, \dots, s_t)$ could be summarized by finite-dimensional sufficient statistics, we could apply the principles in [17] and solve the associated dynamic programming problem to compute the optimal American option price. Moreover, if the stochastic volatility model (1) and (2) could be encompassed within a linear, Gaussian state-space framework, then Kalman-filtering methods would apply to facilitate the exact solution to the pricing problem with partial information. The modeling structure of our problem, however, entails a non-linear, non-Gaussian framework where Kalman-filtering is sub-optimal for posterior inference on the latent process.

The analysis in [58] uses key “summary statistics” of $p(V_t|s_1, \dots, s_t)$, which can be denoted by Q_t , in order to solve the American option pricing problem from a practical perspective.² The work in [58] utilizes an OU process to model log-volatility, and approximated the posterior distribution of the log spot volatility using a Normal distribution. As the dimension of Q_t grows, however, the more computational expense is required to solve the pricing problem. Yet another potential (and practical) approach would be to use the posterior distribution of the average volatility at each time point. The results in [14] demonstrate that the distribution of average volatility would converge to a Normal distribution in accordance with a Central Limit Theorem. In this case, however, average volatility would be used in place of spot volatility to achieve computational gains.

We implement an approach where we also summarize $p(V_t|s_1, \dots, s_t)$ using summary statistics. We address the notion of the volatility risk premium, and use empirical data, along with our proposed pricing methodology, to make inference on this key parameter. We next discuss an assessment of the volatility risk premium in terms of observed and unobserved volatility.

2.4 Risk Quantification

The volatility risk premium, λ , in the risk-neutral pricing model of (1) and (2) is associated with investors’ risk appetite/aversion. Empirical studies have shown that

²The references in [58] also provide additional background on American option valuation with stochastic volatility.

this parameter is typically negative when estimated from options data (see, e.g., [1, 2, 52]). This implies that investors require a premium for taking on risk associated with volatility. As λ becomes more negative, there is an increase in the volatility drift, thereby increasing the likelihood that spot volatility will increase. An increase in volatility results in an increase in option prices for “plain-vanilla” options like put or call options – i.e., the “vega,” the sensitivity of option price with respect to volatility, is positive³ (cf. [29]). One can estimate an optimal value of λ using numerous approaches. In what follows, we make a distributional assumption on observed American option prices, which is similar to what [24] did for European options.

Depending on whether volatility is observed or latent, there are implications for the estimation of λ . First, for a given value of λ , Bellman’s Principle of Optimality (see [3]) states that the option price computed from an assumption of observed volatility will be at least as great as that computed from an assumption of unobserved volatility. Ultimately, this has implications for the optimized estimate of λ computed from the two pricing approaches, which we summarize below in Proposition 2.1.

Proposition 2.1. *Let P_V denote the price estimate obtained when the V_t process is observed, and let P_Q denote the price estimate obtained when V_t is assumed latent and a summary vector Q_t is used to capture information about it. Suppose that the stochastic volatility model parameters, $\theta = (\alpha, \beta, \gamma)$, are fixed. Let λ represent the volatility risk premium, and assume that L observed option contracts C_i , $i = 1, \dots, L$, are distributed independently such that*

$$C_i \sim \text{Normal}[P(\phi_i, \theta, \lambda), \sigma^2], \quad i = 1, \dots, L,$$

where $P(\phi_i, \theta, \lambda)$ is determined from a pricing model using the pricing inputs, ϕ_i , for contract i (e.g., strike, maturity), and the parameters θ and λ . Additionally, σ^2 is the variance, which for simplicity will be assumed constant. Furthermore,

- When volatility is assumed observed, let λ_V^* be the optimal value of λ in the sense that it minimizes the sum-of-squared distances between observed and model-predicted option prices, and
- When volatility is assumed unobserved, let λ_Q^* be the optimal value of λ in the sense that it minimizes the sum-of-squared distances between observed and model-predicted option prices.

Then, for plain-vanilla American-style option contracts, $\lambda_Q^* \leq \lambda_V^*$.

Proof. Let λ be given, and as noted above, let P_V be an estimator of $P(\phi_i, \theta, \lambda)$, the price of an American-style option, assuming both the share price (S_t) and square of volatility (V_t) are observed. According to the Bellman Principle of Optimality, P_V

³The vega of more exotic options (e.g., options on spreads) may not necessarily be positive. See [39] for additional discussion.

is optimal in the sense that it results in the most “optimal policy” (i.e., it produces the optimal stopping rule) for the given value of λ .

As a result of this principle, P_V produces the highest value for the American option price with the given value of λ . Hence, any other price estimator, particularly one that relies on an estimate of volatility rather than the actual observed volatility, will result in a lower option price than P_V .

Let

$$R_i(P, \lambda) = C_i - P(\phi_i, \theta, \lambda),$$

where, $R_i(P, \lambda)$ takes two arguments such that (a) P represents a price estimator of $P(\phi_i, \theta, \lambda)$, and (b) λ represents a value of the market price of volatility risk. Note that from the assumption of the proposition,

$$E[R_i(P, \lambda)] = 0.$$

Recall P_V is the argument of P that represents an estimator of $P(\phi_i, \theta, \lambda)$ when both S_t and V_t are observed. Our objective is to find the optimal value of λ given a price estimator of $P(\phi_i, \theta, \lambda)$. Assume that λ_V^* is the optimal value of λ when P_V is used to estimate $P(\phi_i, \theta, \lambda)$ such that λ_V^* satisfies

$$\lambda_V^* = \arg \min_{\lambda} \sum_{i=1}^L R_i^2(P_V, \lambda).$$

Now let P_Q be the argument of P that represents an estimator of $P(\phi_i, \theta, \lambda)$ when only the S_t process is observed, and the summary vector Q_t is used to capture information about the latent V_t process.

Suppose further that λ_V^* is the value of λ used when P_Q is used as an estimator of $P(\phi_i, \theta, \lambda)$. According to the Bellman Principle of Optimality, $P_Q < P_V$ since P_V results in the highest price (most optimal stopping rule) for a given value of λ . Hence, the P_Q price estimate of $P(\phi_i, \theta, \lambda)$ would be biased low, so the residual $R_i(P_Q, \lambda_V^*)$ would be such that

$$E[R_i(P_Q, \lambda_V^*)] \neq 0,$$

a direct violation of the modeling assumption of the proposition. Moreover, the residual sum-of-squares would not be minimized. Consequently, in order to increase the price estimates of P_Q and make them unbiased to conform with the modeling assumptions stated in the hypothesis, we need to find the value of λ , say λ_Q^* , that minimizes the residual sum-of-squares, i.e.,

$$\lambda_Q^* = \arg \min_{\lambda} \sum_{i=1}^L R_i^2(P_Q, \lambda).$$

In order to find λ_Q^* , we must search to the left of λ_V^* because it is in this range that values exist which will increase the drift rate in (2), thereby increasing the volatility, and hence increasing the American option price based on P_Q . \square

The above proposition formalizes the notion that investors require a larger premium for assuming the risk associated with a *latent* stochastic volatility process relative to an *observed* stochastic volatility process. Specifically, investors require more compensation for bearing risks associated with adverse movements in the market when they make posterior inference on the spot volatility compared to when they are able to observe it directly.

3 American Options and Particle Filters

Particle filter methods for American-style options using the OU model for log-volatility model are explored in [58], and a computing supplement in the R programming language is available in [59]. Particle filters are only relevant to American options in a partial information modeling framework. We present our pricing algorithm using particle filters to summarize $p(V_t|s_1, \dots, s_t)$, which is the posterior filtering distribution of the latent variance process. These filters are based on the sequential importance sampling/bootstrap filter as described in Chap. 1 of [23]. There are several methods to improve the performance of particle filters, such as auxiliary based filters (see [55, 56] for additional details). Indeed, additional adaptive methods like “particle learning” are also available (see [12]). A basic diagram that illustrates the fundamentals of the particle filter cycles appears in Schema 1.

Schema 1: Particle filter posterior/predictive relations

$$p(V_t|s_1, \dots, s_{t-1}) = \int p(V_t|V_{t-1}) p(V_{t-1}|s_1, \dots, s_{t-1}) dV_{t-1} \quad (6)$$



$$p(V_t|s_1, \dots, s_t) \propto p(s_t|V_t) p(V_t|s_1, \dots, s_{t-1}) \quad (7)$$

Equations 6 and 7 in Schema 1 are presented in similar forms in [23] or [45]. These two steps that involve prediction and filtering lie at the core of a particle filtering algorithm. First, (6) calls for sampling from the transition density of the latent process (in this case V_t) to produce predictive draws of the variance process. Second, in (7), these predictive draws are resampled using the likelihood $p(s_t|V_t)$ at the current time point t as weights. These two steps are cycled over time as more data

are collected. (Note that we have emphasized the realized values of the underlying observed process S_t as s_t .) We next discuss how we use filter based statistical summaries in our proposed American option particle filter pricing algorithm.

3.1 Filter Statistics

Generally, it will be impractical to execute a dynamic programming algorithm that fully accommodates the posterior filtering distribution $p(V_t | s_1, \dots, s_t)$. Unless this distribution is parameterized by finite-dimensional sufficient statistics, some type of approximation will need to be made. One could integrate out the latent V_t using the draws from $p(V_t | s_1, \dots, s_t)$, but this would be computationally prohibitive. In a grid-based pricing algorithm, such as the one in [57], parametric approximations (say, based on a specific distribution, or perhaps a mixture of Normals) could be used. A Monte Carlo pricing algorithm, such as the LSM algorithm, could use key summary statistics of the filtering distribution to approximate the price of an American option. Specifically, the filter statistics enter the LSM regression at each exercise/hold decision point as explanatory variables.

Figure 1 presents an illustration of filtering distributions for a price series, assuming that the data follow a square-root mean-reverting stochastic volatility model.⁴ As can be seen, the filtering distribution vary from time point to time point. For example, skewness or kurtosis may be more pronounced in one instance relative to another. Furthermore, multiple modes may exist at some time points but not at others. The types of summary statistics that are relevant may also change from one modeling framework to the next. A determination of the key summary features to use can be gleaned from an analysis of historical price series.

Arguably, measures of center (mean, median) and spread (variance, IQR) will be useful as covariates in the LSM regressions for most types of stochastic volatility models. The square-root mean-reverting stochastic volatility model has skewness built into the transition density, so a measure of skew would also be useful in this case. One could also incorporate additional filter statistics in order to improve accuracy, however, this will entail larger computational costs. In our numerical experiments and empirical analysis, the use of the first three moments of the filtering distribution sufficed. An assessment must be made that practically balances accuracy and practicality for the specific type of stochastic volatility model used.

3.2 Pricing Algorithm

The pricing algorithm below is a fusion of the least-squares Monte Carlo (LSM) approach of [46] and a sequential Monte Carlo based routine as found in [23]. The

⁴Additional details about the data set/data analysis will be available in Sect. 5; the discussion here is only meant to be illustrative.

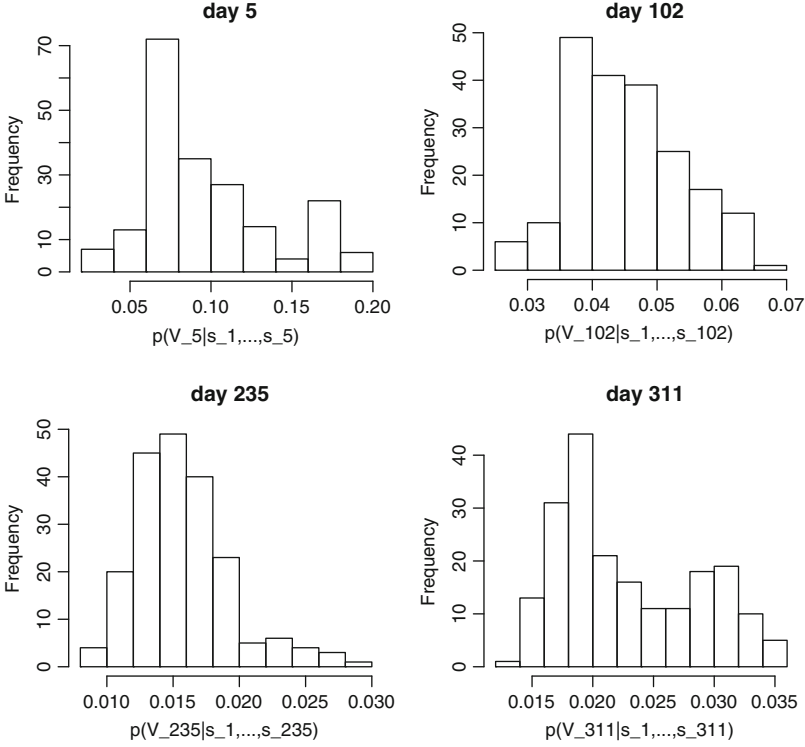


Fig. 1 Illustrative filtering distributions, $p(V_t | s_1, \dots, s_t)$, at various time points

sequential Monte Carlo analysis is implemented directly on the Monte Carlo paths that are simulated in the LSM algorithm. At each time step, we need to compare the higher of the exercise or hold value of the American option. If there are M paths and N time points, the exercise value is specified by a payoff function, $g(S_{i,j}, j)$, and the (approximate) hold value in our latent pricing framework, $H_{i,j}$, $i = 1, \dots, M$, $j = 1, \dots, N$, is given by

$$H_{i,j} = E[\pi(i, \tau_{j+1}) | S_j, Q_j], \quad (8)$$

where, τ_{j+1} , according to the Bellman Principle, is the next optimal time of exercise after time point j . In [46], this is referred to as a “discounted future cash flow” in the path simulation framework.

We present the steps of our pricing approach below, however, it should be noted that substantial computational gains can be achieved by using more efficient versions of the LSM algorithm and/or the sequential Monte Carlo routine.

Routine 3: A particle filter based American option pricer

1. **Path simulation.** For each $i = 1, \dots, M$, and each $j = 1, \dots, N$,
 - Use Routine 2 to get simulated values of the price process $\{s_{i,j}\}$ on a discrete set.
 - At time $j = N$, the expiration date of the option, compute the payoff value using the function $g(S_{i,N}, N)$.
2. **Filter summarization (see [23]).** For each $i = 1, \dots, M$, and each $j = 1, \dots, N - 1$,
 - *Predictive sample:* Use step 3 of Routine 2 and (5) and (6) to simulate m “particles” $\{v_{i,j}^k\}$ from $p(V_t | s_1, \dots, s_{t-1})$, where $k = 1, \dots, m$, using $p(V_{i,j} | V_{i,j-1})$. (See (6) in Schema 1.)
 - *Weights:* Compute “weights,” $w_{i,j}^k$, by evaluating $p(s_{i,j} | s_{i,j-1}, v_{i,j}^k)$ for $k = 1, \dots, m$. (See (4).)
 - *Posterior sample:* Resample from the predictive particles, $\{v_{i,j}^k\}$, with weights proportional to $p(s_{i,j} | v_{i,j}^k)$ to obtain draws from $p(V_t | s_1, \dots, s_t)$, which will serve as input to the next iteration of the prediction step above. (See (7) in Schema 1.)
 - Summarize $p(V_{i,j} | s_{i,1}, \dots, s_{i,j})$ in a “summary vector” $Q_{i,j}$ that captures measures of center, spread, and skew, and any additional key distributional features.
3. **Decision step.** At each time point $j = N - 1, \dots, 1$,
 - Compute the exercise value by evaluating the payoff function, $g(S_{i,j}, j)$, along each path $i = 1, \dots, M$.
 - *LSM sub-step:* Compute the hold value by regressing the first instance of discounted future cash flows from the M paths on basis functions of the M values of the observed series at time point j , $\{s_{1,j}, \dots, s_{M,j}\}$ as well as of the summary vector $\{q_{1,j}, \dots, q_{M,j}\}$. (See [46] for full LSM algorithm.)
 - Evaluate the higher of the exercise or hold value, flagging the instances along each of the M paths where the exercise value is higher, thus creating an instance of “future cash flow” for previous steps in the LSM-based dynamic programming algorithm.
4. **Price estimate.** For each path $i = 1, \dots, M$,
 - Discount the cash flow from the first instance of exercise to the valuation time, t , and average to get a Monte Carlo based American option price, and
 - Compute approximate standard errors using the Normal approximation to the Monte Carlo average.

Remark 3.1. Concerning the filter summarization step above, at time $t = 0$ the initial sample of particles from $p(V_0|s_0)$ or $p(V_0)$ can be drawn from the stationary distribution of V_t under the statistical measure, which is Gamma $\left(\frac{2\alpha\beta}{\gamma^2}, \frac{2\alpha}{\gamma^2}\right)$. Alternative choices for this initial distribution can also be implemented.

Remark 3.2. We have used lower case $s_{i,j}$ and $v_{i,j}$ to denote simulated realized values.

A more detailed version of the above algorithm appears in [58], and there is also a grid-based version of the algorithm in the aforementioned work as well. As noted above there are different variants of the particle filtering algorithm and the LSM algorithm. The LSM algorithm can be used with all paths or only “in-the-money” paths for improved efficiency. It should also be noted that the Monte Carlo standard errors are approximate because they rely on independence of the price estimates along each path. There is dependence, however, introduced by the regression function that is used to estimate the hold value at each decision step. Moreover, the work in [62] discusses how to compute a “dual price” for American options where a lower bound price, from say the LSM algorithm, is combined with a dual upper bound price.

4 Benchmark Analysis

We now present the results of a small benchmark analysis where we compare the results of our pricing methodology in the “limited information” setting to one where the full information set is available. A portion of the work in [58] shows that stochastic volatility actually matters, and that pricing estimates are markedly different if sub-standard estimates of volatility are used in a latent stochastic volatility framework. Concerning the OU model for log-volatility, the work in [58] demonstrates that a sequential Monte Carlo based pricing result comes within a standard error (sometimes less) of the result from a pricing approach that assumes volatility is observable.

There are a number of ways to numerically assess the accuracy of Routine 3. One could compare the true filtering distribution $p(V_t|s_1, \dots, s_t)$ with a finite-dimensional approximation that is parameterized by the summary statistic vector Q_t . For example, if Q_t summarizes the center, spread, and skew, a comparison could be made to a skew-normal or a two-component Normal mixture that would capture skew. We could also compare the distributions $p(S_t|s_1, \dots, s_{t-1})$ with $p(S_t|Q_t)$, noting that Q_t contains summaries that encapsulates $\{s_1, \dots, s_{t-1}\}$. Although computationally intensive, these comparisons could be made using standard distance measures such as the Kullback–Leibler (KL) distance, or on test-statistics based on these measures.

We compare the obtained prices obtained from Routine 3 (particle filter pricer) for the limited observation case to those produced by the full observation case under the following scenarios: (a) high mean–reversion/low volatility–of–volatility, and low mean–reversion/high volatility–of–volatility, (b) degrees of moneyness, and (c) length of maturity. The full–state algorithm is the standard LSM algorithm assuming that both processes, S_t and V_t , are observed state variables. Therefore, it is straightforward to price American options in the full information case since one can simulate paths from the given model (e.g., (1) and (2), and then use these simulated values as predictors in the LSM regression steps.

We fix the number of Monte Carlo paths $M = 10,000$ and the number of particles in the sequential Monte Carlo algorithm to be $m = 200$ in our numerical experiments.⁵ Additionally, we fix the values of the risk-free rate $r = 0.01$, the correlation $\rho = -0.25$, and the volatility risk premium $\lambda = -5.05$. The prior distribution on the initial variance, $p(V_0)$, plays a key role in the option price results as they are intricate convexity effects on the price due to volatility. We use the stationary distribution of the square–root mean–reverting stochastic volatility model, which is distributed as a Gamma as described in Remark 3.1 under Routine 3 above.

Additionally, the experiments comply with the positivity constraint, $\gamma^2 < 2\alpha\beta$, which is integral to the square–root mean–reverting stochastic volatility model utilized below. We parameterize the summary vector Q_t by the first three moments of the filtering distribution, $p(V_t|s_1, \dots, s_t)$, to capture measures of center, spread, and skew. We also use the first two Laguerre polynomials evaluated at the share price and the three components of the summary vector as well as all possible cross–terms associated with these in the LSM algorithm. In the following numerical experiments, we price American–style put options, which have a payoff function equal to

$$g(S_t, t) = \max(K - S_t, 0),$$

where K is the strike price and S_t is the observed underlying at time t .

Table 1 illustrates the differences between American option pricing results using a limited information approach (Routine 3) and a full–state information approach. In the case of low mean–reversion and high volatility–of–volatility (row 1), there may be an argument that the limited information result is noticeably less than the

Table 1 Differences in limited information vs. full information pricing approaches in terms of mean–reversion and volatility–of–volatility. The initial settings used to simulate the paths are $S_0 = \$48$, $V_0 = \beta$, $K = \$50$, and $T = 20$ days. Monte Carlo standard errors are in parentheses

(α, β, γ)	Limited observation	Full observation
(0.1, 7.1, 0.9)	15.17 (0.076)	15.31 (0.075)
(50, 0.05, 0.03)	2.04 (6.92e–03)	2.06 (6.95e–03)

⁵Although these are small simulation sample sizes relative to large–scale Monte Carlo experiments, they are suitable for our illustrative purposes.

full information result. This might be expected since volatility is indeed a dominant stochastic factor in this setting. The difference, however, is still within a reasonable margin of the Monte Carlo error. In the case of high mean–reversion and low volatility–of–volatility (row 2), the difference between the limited information and the full information cases is negligible. If the limited and full information option price results are similar, then this is evidence that the proposed partial observation pricing algorithm (Routine 3) is effective in capturing information about the latent volatility process.

Table 2 measures the differences between the limited information and full information results for varying degrees of moneyness (out–of–the–money, at–the–money, and in–the–money). Indeed, in all cases, the pricing results are well within the Monte Carlo standard error of each other. Table 3 also shows small differences between the limited and full observation pricing methods. Perhaps a minor argument can be made that for short-dated American options, the sequential Monte Carlo algorithm may not gather enough information from the observed data before the option expires in order to accurately “learn” about the latent volatility. Thus, the pricing result is slightly less than the full information instance. The difference, nonetheless, is too small to make a definitive conclusion.

The experiments above are on a small scale, however, they demonstrate that a sequential Monte Carlo based approach is very effective for pricing American–style options in a limited information setting. They are also robust with respect to the choices we made for the components of Q_t and the number of Laguerre basis functions we used in the LSM algorithm. Adding more components to the summary vector Q_t , however, will result in improved inference on the latent volatility process. Ultimately, although combining posterior inference on the latent quantity with the optimal stopping problem is computationally intensive, the benefits in terms of risk management decisions could far outweigh the costs. We now turn to an empirical exercise where inference is made for the market price of volatility risk, λ .

Table 2 Differences in limited information vs. full information pricing approaches in terms of moneyness. The initial settings are $S_0 = 48$, $V_0 = \beta$, $\alpha = 0.9$, $\beta = 0.3$, $\gamma = 0.1$, $K = \$50$, and $T = 20$ days. Monte Carlo standard errors are in parentheses

S_0 (\$)	Limited observation	Full observation
45	5.46 (0.018)	5.45 (0.018)
50	1.04 (0.012)	1.04 (0.012)
55	0.010 (8.26e–04)	0.013 (1.06e–03)

Table 3 Differences in limited information vs. full information pricing approaches in terms of maturity length. The initial settings are $S_0 = \$49$, $V_0 = \beta$, $\alpha = 0.9$, $\beta = 0.3$, $\gamma = 0.1$, and $K = \$50$. Monte Carlo standard errors are in parentheses

T (days)	Limited observation	Full observation
5	1.16 (7.93e–03)	1.19 (8.72e–03)
25	1.87 (0.016)	1.90 (0.017)
50	2.58 (0.022)	2.62 (0.024)

5 Application to Index Options

One convenient result of developing an option pricing algorithm is the ability to make statistical inference by extracting information from observed market prices. Several studies have implemented empirical analysis, mostly for European-style options, in order to estimate risk-neutral model parameters. Some important examples include [13, 24, 36, 52, 75], and some of these studies specifically focus on stochastic volatility models as well as jump processes. Estimation of model parameters under the statistical measure requires only information on the observed share prices. On the other hand, estimation of model parameters under the risk-neutral measure requires data on the share prices and market-observed option prices. In a European option pricing framework, the estimation of risk-neutral parameters, while challenging, is computationally feasible as is demonstrated for the square-root mean-reverting model in [52] and [24].

Due to the early-exercise feature of American-style options, joint estimation using both share and option prices is far more computationally challenging relative to the European option pricing framework. In an observed stochastic volatility setting, the computational problem is feasible to solve since standard dynamic programming methods are accessible. In the limited information setting that is of interest in this analysis, the combination of sequential Monte Carlo filtering and dynamic programming requires significant computing resources. There are many sophisticated enhancements of particle filtering to address inference in a sequential framework, as is inherently the case for American options, and one good example is, among others, [44].

In the analysis that follows, we follow the general estimation methodology as implemented in [58, 59], however, we adapt it to the square-root mean-reverting model. As is the case for the benchmark analysis above, we use the first three moments of the volatility filtering distribution to parameterize the summary vector, and we use the first two Laguerre polynomials in the share price and the components of the summary vector, as well as all cross-terms, in our application of the LSM algorithm. Other choices are available and these may depend on the type of stochastic volatility model employed for analysis. We “split” the estimation of model parameters in two parts: one under the statistical measure and the other under the risk-neutral measure. Additionally, motivated by Proposition 2.1, we use market data and report observations on the perception of volatility risk premium for American-style options.

5.1 Data Description

Inference on model parameters under a risk-neutral framework requires data on both the underlying price series, S_t , and the option contracts, C_t . We use daily closing prices of the “NYSE Arca Oil Index” as the underlying price series

that we obtain from Bloomberg[®]. We obtain data on the index from Jan. 1st, 2003 to Mar. 31st, 2004, which we use as the historical period, for estimation of the stochastic volatility model parameters under the statistical or “real-world” measure. As noted on the NYSE Euronext (New York Stock Exchange) website, this index, symbolized by “XOI,” is a “price-weighted index designed to measure the performance of the oil industry through changes in the prices of a cross section of widely-held corporations involved in the exploration, production, and development of petroleum.” Additionally, the index has a “benchmark value” of 125.00, which was established on Aug. 27th, 1984.⁶

We also obtain daily closing prices on American-style put options on the NYSE Arca Oil Index from the NYSE Euronext website. These data span the period Apr. 1st, 2004 to Jun. 21st, 2004, which we call the “valuation period.” We also acquire the corresponding data on the underlying NYSE Arca Oil Index during this period. We analyze American put options written on this index since American puts are canonical examples of early-exercise financial derivatives. This present empirical analysis, which treats index options, could be contrasted to the work in [58] where a similar analysis is done using American put options on equities, and option contracts on three equities are selected for that analysis.

A plot of the NYSE Arca Oil Index appears in Fig. 2 for the period Jan. 1st, 2003 to Mar. 31st, 2004. Recall that we use this data period for estimation of model




Fig. 2 Historical data series for NYSE Arca Oil Index (XOI) during the period Jan. 1st, 2003–Mar. 31st, 2004, or first-quarter of 2003 to first-quarter of 2004, inclusive

⁶See <http://www.nyse.com> for additional details on the NYSE Euronext index options.

parameters under the statistical measure. Generally, over the year 2003 and into 2004, the overall trend for the oil industry is an increasing one. Note that we use data for the four quarters of 1 year and one quarter into the next to avoid potential, abrupt “year-end” effects, and because our option sample spans the second quarter of 2004.

A summary of the XOJ Index for the historical and valuation periods appears in Table 4. Key measures of the raw data set, including measures of center and spread, are reported. The essence of our empirical analysis will be to combine model parameter estimation results from these data with the option data. Note that while a joint analysis would be an enhancement to this empirical exercise, we separate out the analysis into the historical period and valuation period for computational purposes.

A plot of the data utilized for the valuation analysis appears in Fig. 3. The top left plot shows the underlying XOJ Index over the second quarter of 2004. The remaining three plots each show the price series for three American put options on the XOJ Index, differing in terms of strike and maturity date. As can be seen, the option prices tend to exhibit some “stickiness” during portions of the valuation period (i.e., the price remains unchanged for a few days). We selected these options, however, because they are more liquid in that they are typically “near-the-money” and have a maturity between 1 and 6 weeks. The “open interest” of these options, an approximate gauge of their trading activity, ranged between 6 and 1,429 contracts, with an average of about 201 contracts, over this data period. According to the NYSE Euronext website, the notional value for each contract equals \$100 multiplied by the underlying index value.

Table 5 presents summary information on the three option series. Panel A presents statistical summary information on the option series, and panel B describes the “physical features” of the option contracts. Clearly, there are potentially countless random sources that influence option prices. In this analysis, we only focus on the volatility of the underlying observed process, the XOJ Index, and make inference on the model–predicted option price assuming that stochastic volatility is latent. We now discuss our analysis of these data and report inferential results along with any notable implications.

Table 4 Numerical summaries for the XOJ Index over (a) the historical period: Jan. 1st, 2003 to Mar. 31st, 2004, and (b) the valuation period: Apr. 1st, 2004 to Jun. 21st, 2004

Measure	Historical data	Valuation data
Min.	411.10	586.10
First-quartile	446.90	602.20
Median	481.00	607.80
Mean	489.70	607.90
Std. dev.	53.81	10.988
Third-quartile	515.10	613.70
Max.	606.60	634.60




Fig. 3 American put option series for NYSE Arca Oil Index (XOI) for various sub-periods during second-quarter of 2004

Table 5 *Panel A:* Price summary for American put option series on the XOI index over the period: Apr. 1st, 2004–Jun. 21st, 2004. *Panel B:* Summary information on the option data set characteristics

Measure	Series 1	Series 2	Series 3
<i>Panel A:</i>			
Min.	3.50	3.90	1.15
First-quartile	7.10	9.50	0.25
Median	12.50	14.30	18.30
Mean	12.26	14.14	20.78
Std. dev.	6.86	5.70	9.17
Third-quartile	14.70	17.80	0.75
Max.	24.00	21.70	36.90
<i>Panel B:</i>			
No. of options (L)	33	31	29
Strike price	600	610	620
Maturity date	05/24/2004	06/21/2004	06/21/2004

5.2 Parameter Estimation

We implement parameter estimation in two steps: first under the statistical measure, and second under the risk-neutral measure. Our discussion in this sub-section focuses on estimation of model parameters under the statistical measure. Note that under the statistical measure, (1) and (2) will be revised since (a) the risk-free rate r will be replaced by the physical drift rate, denoted by, say, μ , and (b) λ , the volatility risk premium, is set to 0 (or does not appear in the expression for the stochastic process governing volatility). We will not report results on the drift rate μ as it is not integral to the pricing problem. We use the 3-month LIBOR rate for r , obtained from Bloomberg[®], during the second quarter of 2004, which is close to 0.01.

Since we assume that volatility is latent, the likelihood function is not available in closed-form. This poses additional computational challenges for optimization of the likelihood (or log-likelihood) function. In fact, particle filtering methods, as demonstrated in [42] and [43], are typically used to create a Monte Carlo approximation to the likelihood function. This Monte Carlo based likelihood is then optimized using either a maximum likelihood based approach, or a Markov chain Monte Carlo (MCMC) approach if prior distributions are specified on the model parameters.

We utilize a simulated annealing (SA) algorithm, which is a stochastic search algorithm, in order to optimize our Monte Carlo based likelihood function for the model in (1) and (2). The SA algorithm is well-suited to search in a high-dimensional parameter space, and will be useful in finding approximations to the optimal values of $(\rho, \alpha, \beta, \gamma)$. The SA optimization approach is discussed at length in [61], and has been studied by [63] in the context of state-space models. It is essentially an MCMC type algorithm, however, prior distributions are not specified. Typically, a proposal distribution is required to execute the stochastic search. We experimented with both a multivariate Normal and a multivariate Student- t distribution as proposal candidates, however, we did not find any appreciable differences in the output, thus we used a Normal proposal distribution. Additionally, there is a “temperature” parameter that is altered according to a “cooling schedule,” which enables the algorithm to converge to the optimized parameter estimates. In Routine 4, we present a summary of the steps of the SA algorithm based on [61].

Routine 4: Simulated annealing procedure

1. Initialization.

- Start with an initial (current) values, and set an initial value for the cooling parameter denoted by C (e.g., 100).
- Let $i = 1, \dots, B$, $t = 1, \dots, T$, and $k = 1, \dots, m$ be the SA, data series, and particle filter loop indices, respectively.

2. *Parameter search.* For each $i = 1, \dots, B$,

- Sample a candidate parameter value using a suitable “proposal distribution” (e.g., Normal or Student- t).
- Cycle through the first three sub-steps of step 2 (“Filter summarization”) of Routine 3, and compute and store the weights, $w_{t,k}$, for $t = 1, \dots, N$, $k = 1, \dots, m$, for the candidate and current parameter values.
- Approximate the log-likelihood function by computing

$$\text{LLik} = \sum_{t=1}^T \frac{1}{m} \sum_{k=1}^m \ln(w_{t,k})$$

for each of the candidate and current parameter values.

- Compute $D = \exp \{[\text{LLik}(\text{candidate}) - \text{LLik}(\text{current})] / C\}$
- Sample $u \sim \text{Unif}(0, 1)$. If $u < D$ accept candidate parameter, else reject it.
- If candidate parameter accepted, assign it as the current parameter.
- Decrease C by a factor between 0 and 1 (e.g., 0.90).

3. *Results.*

- Output the results of the SA search and check plots for adequate “mixing” of the SA chain.
- Utilize about last 5% of draws from the SA chain for subsequent analysis.

Remark 5.1. Concerning the proposal distribution, if it is symmetric, the mean could be set at the current point, and the covariance matrix should be parameterized to ensure adequate coverage of the parameter space.

Remark 5.2. In the event that a given draw from the parameter space does not satisfy the constraints of the square-root mean-reverting model, the draw ought to be automatically rejected.

Remark 5.3. To achieve additional computational efficiency, the parameters can be reparameterized and sampled in blocks in the SA search step.

The SA algorithm, if executed appropriately, converges to the point that maximizes the likelihood function. We report the XOI Index results for the SA algorithm in Table 6 with 25,000 iterations, and we used the last 1,000 iterations to make inference.⁷

⁷We used the last 1,000 draws, as opposed to the last accepted draw, to compute simulation-based distributional summaries. The results are also robust across longer iteration lengths.

Table 6 Optimized posterior estimates of the model parameters under the statistical measure. Approximate 95% posterior credible intervals are also reported

Parameter	Mean	95% credible interval
ρ	-0.799	(-0.874, -0.729)
α	0.591	(0.363, 0.974)
β	0.0479	(0.0287, 0.0693)
γ	0.166	(0.0789, 0.243)

The results in Table 6 point out some well-known observations about stochastic volatility models. First, we note that the correlation parameter is negative, and the approximate 95% credible set does not include 0. The correlation is between the Brownian motions that shock the underlying and the volatility processes. This confirms what is known as the “leverage effect” in equities and indices, and several earlier studies document this finding. Additionally, we report estimates (and approximate credible sets) of the mean–reversion rate (α), the mean–reversion level (β), and the volatility of the process V_t (γ) under the statistical measure.

The output in Fig. 1 uses these parameter estimates to approximate the illustrated posterior distributions of the variance process, V_t , at select points in the historical and valuation periods. The summary vector, Q_t , used in the LSM algorithm to price American-style options, are based on distributions similar to these. We next discuss the estimation steps under the risk–neutral measure, conditional on the estimation results from the statistical measure, and what it entails for the market price of volatility risk, λ .

5.3 Volatility Risk Premium

The volatility risk premium, λ , reveals information that can help assess the risk profile of holders of option contracts. Inspection of (2) reveals that as λ becomes more negative, the drift rate in the variance process increases, which will tend to cause the spot variance, and hence the spot volatility, to increase. As is known in the literature, this will generally result in an increase in option prices. Therefore, as λ gets more negative, option prices increase, which signals that investors require additional compensation for bearing risks associated with stochastic volatility. An interesting question to pose is whether there are differences in investors’ risk preferences across time or in terms of observability.

The fact that volatility is not a traded asset entails that there is no unique specification for λ . We assume that it is a constant in this analysis, but it could be time–varying, or it could even follow its own, separate stochastic process. Assuming that we have observed prices on L American put option contracts, $C_i, i = 1, \dots, L$, we can minimize the residual sum–of–squared errors between observed and model–predicted option prices. As in Sect. 2.4 above, we can denote by P_V and

P_Q the pricing results based on the observed and unobserved cases, respectively. Consequently, the residual, or the distance between observed and model-predicted option price, could be written as

$$R_i(P_V, \lambda) = C_i - P_V(\phi_i, \theta, \lambda), \text{ and} \quad (9)$$

$$R_i(P_Q, \lambda) = C_i - P_Q(\phi_i, \theta, \lambda), \quad (10)$$

for the observed and unobserved cases respectively. Thus, the corresponding optimized estimates of λ for the observed and unobserved cases would respectively minimize the residual sum-of-squared errors – i.e.,

$$\lambda_V^* = \arg \min_{\lambda} \sum_{i=1}^L R_i^2(P_V, \lambda), \text{ and} \quad (11)$$

$$\lambda_Q^* = \arg \min_{\lambda} \sum_{i=1}^L R_i^2(P_Q, \lambda), \quad (12)$$

Numerous studies have addressed the volatility risk premium, and a small sample of some include [1, 2, 22, 64, 74]. The work in [53], while contributing to the stochastic volatility option pricing literature, assumes a zero volatility risk premium as well as no leverage effects. It is common in many studies to assume that $\lambda = 0$, and in fact, this is usually referred to as the “minimal martingale measure” (see [50]). Furthermore, the additional complexity of constructing posterior distributions of the spot volatility accounting for the market price of volatility risk is rarely ever addressed.

To solve the optimization in (11) and (12), we use the setup in [58] whereby a grid-search approach is adopted. A key objective is to make statistical inference on the volatility risk premium based on the results of the optimization in (11) and (12). The details of our grid-search methodology are summarized in the following routine.

Routine 5: Volatility risk optimization

1. Initialization.

- Obtain prices on L market observed American option contracts across various strikes and maturities.
- Create a grid of values, $g = 1, \dots, G$, for the market price of volatility risk λ .

2. Grid search. For each $g = 1, \dots, G$,

- Use Routine 3 to compute the model-predicted American option price for each of the L option contracts.

- Compute the sum-of-squared residuals between model-predicted and observed option prices for the L option contracts and denote this by R_g .

3. Result.

- Output the value of λ that corresponds to the minimum of the set $\{R_1, \dots, R_G\}$.

Remark 5.4. The above grid-based optimization is for finding λ_O^* , the optimal value of the volatility risk premium in the partial information (unobserved volatility) case. For the full information case, one would use the standard LSM algorithm with S_t and V_t as state variables in lieu of Routine 3 in step 2 above in order to find λ_V^* .

It can be shown that the resulting optimized quantity is equivalent to the maximum likelihood estimate (MLE) of the parameter λ , or the least-squares estimate from a non-linear regression perspective, using the Gaussian modeling framework for option prices outlined in Proposition 2.1. Moreover, as outlined in [66], a vague prior specification results in a convenient posterior distribution whose maximum (or mode) is equivalent to the MLE. We approximate confidence intervals for the optimized estimates of λ using finite differences.

In contrast to the work in [58], we analyze the volatility risk premium from two different perspectives. First, we investigate whether or not the volatility risk premium for American-style option contracts differ across time. Second, we examine whether there is an observability effect that is recognized from the optimized volatility risk premium. In order to implement this small study, we use the option contract data, as well as the underlying price data, on the XOI Index described in Sect. 5.1. Regarding the first question, we follow an option series for the month prior to its expiration date, and also for the month when it expires, using as much data as is available. Concerning the second question, we use the standard LSM algorithm and the particle filter based Routine 3 to price the options under the assumption of observed volatility (“full observation”) and latent volatility (“limited observation”), respectively. Our results are reported in Table 7.

The results reveal some intriguing findings about time and observability effects on the volatility risk premium. First, we see that the estimates of λ are negative, and their approximate 95% confidence intervals do not span zero. This confirms the extant empirical findings that the volatility risk premium is negative, which implies that investors require compensation for bearing risk due to adverse movements in volatility. Second, this small empirical exercise confirms what is claimed and proved in Proposition 2.1 concerning λ_O^* being less than or equal to λ_V^* . Note that the magnitude of the volatility risk premium is arguably large, and this is most likely due to the underlying being an index as opposed to an equity.

Table 7 Optimized estimates of the volatility risk premium, along with approximate 95% confidence/credible intervals, for the limited and full observation cases under the risk-neutral measure. Inferential results are conditional on the optimized model parameter estimates for ρ , α , β , and γ under the statistical measure. Comparisons are made across time as well. These estimates are based on the 93 options described in Table 5 of Sect. 5.1

Expiration month?	λ_Q^*	λ_V^*
No	−67.895 (−69.542, −66.248)	−64.789 (−66.910, −62.668)
Yes	−122.517 (−123.845, −121.189)	−118.737 (−119.995, −117.479)

Next, we observe that there appears to be a statistically significant time effect associated with the volatility risk premium. It is significantly *more* negative during the month in which the option contract matures. This is true for both the observed and latent pricing frameworks. Hence, this signals that holders of American-style XOI Index options become more “risk conscious” and require added compensation as the expiration date nears. This finding is also corroborated by work in an earlier analysis by [6] where it is empirically shown that owners of American options exercise closer to the expiration date. In other words, the “market exercise boundary” is not necessarily the optimal exercise boundary. Thus, as λ becomes more negative as the option nears expiration, this may signal the exercise activity, and increased “agitation,” that occurs for American options.

Finally, there is some moderate difference between λ_V^* and λ_Q^* , the optimized volatility risk premium for the observed and unobserved frameworks, respectively. The difference is only statistically significant during the month of expiration. In other words, during the month prior to expiration, whether volatility is observable or not only has a marginal impact on the volatility risk premium, and the impact is not statistically significant. On the other hand, during the month in which the option contract expires, there seems to be an observability effect associated with the volatility risk premium. Specifically, in the limited observation framework where volatility is latent, investors require additional compensation relative to the full information setting where volatility is observed. Thus, this empirical finding confirms what may be intuitive, however, it also marks the time frame where the difference due to observability effects is prominent.

6 Concluding Remarks

We have discussed a pricing methodology for American-style options in a stochastic volatility framework, where it is assumed that volatility is a latent process. Our approach is Monte Carlo based, and it is a fusion of the LSM algorithm of [46] and the particle filter methodology described in [23]. We extend the work in [58]

by working with the more popular square-root mean-reverting model of [33] for stochastic volatility. The contribution of our methodology is to combine the dynamic programming algorithm, associated with the optimal stopping problem for early-exercise options, with sequential posterior inference on the spot volatility (or spot variance). Although our limited observation approach is computationally intensive, it results in prices that mostly are within negligible statistical error of the full observation benchmark. Furthermore, the speed and practicality of our approach can be enhanced with the use of parallel computing architecture.

Apart from providing a framework for pricing American options in a limited observation setting, we offer some observations concerning the volatility risk premium. The premium associated with volatility risk is one of the key risk-neutral parameters, and its estimate can reveal insights about investors' behaviors. We estimate all model parameters, as well as the volatility risk premium, using share price and American put options on the NYSE Arca Oil Index (XOI). We find negative estimates for the volatility risk premium, which signifies that investors' require compensation for bearing risk associated with stochastic volatility. We also find time and observability effects on the volatility risk premium, and we find that investors are more conscious about volatility risk as the expiration date of an American option approaches.

There are several avenues for potential refinements of this work. First, a joint analysis under the statistical and risk-neutral measure, although computationally demanding, would improve parameter estimation. Furthermore, the value of information could be measured by collecting data at varying frequencies and assessing the effect on American option prices. Next, more elaborate specifications of the volatility risk premium, especially in light of the aforementioned time and observability effects, could enhance inference, and ultimately reveal some behavioral insights concerning holders of American options. Our data set is slightly dated, and the behavior of option prices, particularly in the recent volatile market environment, could also be studied. One could more appreciate the benefits of a sophisticated inferential framework for spot volatility when unusual market conditions prevail. The methodology in this analysis could be used for better understanding certain dynamics in financial markets, particularly under volatile conditions.

Acknowledgements The author gratefully acknowledges the suggestions of the Editors and two anonymous referees. Their insights have significantly improved the content of this work.

References

1. Bakshi, G., Kapadia, N.: Delta-hedged gains and the negative market volatility risk premium. *The Review of Financial Studies* **16**(2), 527–566 (2003)
2. Bakshi, G., Kapadia, N.: Volatility risk premiums embedded in individual equity options: Some new insights. *The Journal of Derivatives* **Fall**, 45–54 (2003)
3. Bellman, R.: *An Introduction to the Theory of Dynamic Programming*. Rand Corporation (1953)

4. Black, F., Scholes, M.: The pricing of options and corporate liabilities. *Journal of Political Economy* **81**, 637–659 (1973)
5. Brennan, M.J., Schwartz, E.S.: The valuation of American put options. *The Journal of Finance* **32**(2), 449–462 (1977)
6. Broadie, M., Detemple, J., Ghysels, E., Torres, O.: American options with stochastic dividends and volatility: A nonparametric investigation. *Journal of Econometrics* **94**, 53–92 (2000)
7. Broadie, M., Glasserman, P.: Pricing American-style securities using simulation. *Journal of Economic Dynamics and Control* **21**, 1323–1352 (1997)
8. Broadie, M., Kaya, O.: Exact simulation of stochastic volatility and other affine jump diffusion processes. *Operations Research* **54**(2), 217–231 (2006)
9. Carr, P., Geman, H., Madan, D.B.: Pricing and hedging in incomplete markets. *Journal of Financial Economics* **62**, 131–167 (2001)
10. Carr, P., Jarrow, R., Myneni, R.: Alternative characterizations of American put options. *Mathematical Finance* **2**(2), 87–106 (1992)
11. Carrière, J.: Valuation of the early-exercise price for derivative securities using simulations and splines. *Insurance: Mathematics and Economics* **19**, 19–30 (1996)
12. Carvalho, C.M., Johannes, M.S., Lopes, H.F., Polson, N.G.: Particle learning and smoothing. *Statistical Science* **25**(1), 88–106 (2010)
13. Chernov, M., Ghysels, E.: A study towards a unified approach to the joint estimation of objective and risk neutral measures for the purpose of options valuation. *Journal of Financial Economics* **56**, 407–458 (2000)
14. Chopin, N.: Central Limit Theorem for sequential Monte Carlo methods and its applications to Bayesian inference. *Annals of Statistics* **32**(6), 2385–2411 (2004)
15. Clarke, N., Parrott, K.: Multigrid for American option pricing with stochastic volatility. *Applied Mathematical Finance* **6**, 177–195 (1999)
16. Cox, J., Ingersoll, J., Ross, S.: A theory of the term structure of interest rates. *Econometrica* **53**, 385–407 (1985)
17. DeGroot, M.H.: *Optimal Statistical Decisions*. McGraw-Hill, New York (1970)
18. Dempster, M., Richards, D.: Pricing American options fitting the smile. *Mathematical Finance* **10**(2), 157–177 (2000)
19. Desai, R., Lele, T., Viens, F.: A Monte Carlo method for portfolio optimization under partially observed stochastic volatility. In: 2003 IEEE International Conference on Computational Intelligence for Financial Engineering: Proceedings. IEEE, Hong Kong (2003)
20. Detemple, J., Osakwe, C.: The valuation of volatility options. *European Finance Review* **4**, 21–50 (2000)
21. Detemple, J., Tian, W.: The valuation of American options for a class of diffusion processes. *Management Science* **48**(7), 917–937 (2002)
22. Doran, J.S., Ronn, E.I.: On the market price of volatility risk. Tech. rep., University of Texas at Austin, Austin (2003)
23. Doucet, A., de Freitas, N., Gordon, N.: *Sequential Monte Carlo Methods in Practice*, first edn. Springer, New York (2001)
24. Eraker, B.: Do stock prices and volatility jump? Reconciling evidence from spot and option prices. *The Journal of Finance* **LIX**(3), 1367–1404 (2004)
25. Finucane, T.J., Tomas, M.J.: American stochastic volatility call option pricing: A lattice based approach. *Review of Derivatives Research* **1**, 183–201 (1997)
26. Fouque, J.P., Papanicolaou, G., Sircar, K.R.: *Derivatives in Financial Markets with Stochastic Volatility*. Cambridge University Press, Cambridge (2000)
27. de Freitas, N., Andrieu, C., Højen-Sørensen, P., Niranjana, M., Gee, A.: Sequential Monte Carlo methods for neural networks. In: *Sequential Monte Carlo Methods in Practice*, chap. 17, pp. 359–379. Springer (2001)
28. Geske, R., Johnson, H.: The American put option valued analytically. *The Journal of Finance* **39**(5), 1511–1524 (1984)
29. Glasserman, P.: *Monte Carlo Methods in Financial Engineering*. Springer, New York (2004)

30. Guan, L.K., Guo, X.: Pricing American options with stochastic volatility: Evidence from S&P 500 futures options. *The Journal of Futures Markets* **20**(7), 625–659 (2000)
31. Harrison, J.M., Pliska, S.R.: Martingales and stochastic integrals in the theory of continuous trading. *Stochastic Processes and their Applications* **11**, 215–260 (1981)
32. Heath, D., Platen, E., Schweizer, M.: A comparison of two quadratic approaches to hedging in incomplete markets. *Mathematical Finance* **11**(4), 385–413 (2001)
33. Heston, S.: A closed-form solution for options with stochastic volatility with applications to bond and currency options. *The Review of Financial Studies* **6**(2), 327–343 (1993)
34. Hobson, D., Rogers, L.: Complete models with stochastic volatility. *Mathematical Finance* **8**, 27–48 (1998)
35. Hull, J., White, A.: The pricing of options on assets with stochastic volatility. *The Journal of Finance* **42**(2), 281–300 (1987)
36. Jacquier, E., Polson, N.G., Rossi, P.E.: Bayesian analysis of stochastic volatility models with fat tails and correlated errors. *Journal of Econometrics* **122**, 185–212 (2004)
37. Jasra, A., Moral, P.D.: Sequential Monte Carlo for option pricing. *Stochastic Analysis and Applications*, to appear, (2010)
38. Johnson, N.L., Kotz, S., Balakrishnan, N.: Continuous Univariate Distributions: Volume 2, second edn. John Wiley & Sons, Inc., New York (1995)
39. Kapadia, N.: Negative vega? Understanding options on spreads. *Journal of Alternative Investments* **Spring** (1999)
40. Karatzas, I., Shreve, S.E.: Brownian Motion and Stochastic Calculus, second edn. Springer, New York (1991)
41. Karatzas, I., Shreve, S.E.: Methods of Mathematical Finance. Springer, New York (1998)
42. Kitagawa, G.: Monte Carlo filter and smoother for non-Gaussian non-linear state space models. *Journal of Computational and Graphical Statistics* **5**(1), 1–25 (1996)
43. Kitagawa, G., Sato, S.: Monte Carlo smoothing and self-organizing state-space model. In: Sequential Monte Carlo Methods in Practice, pp. 177–196. Springer (2001)
44. Liu, J., West, M.: Combined parameter and state estimation in simulation-based filtering. In: Sequential Monte Carlo Methods in Practice, chap. 10, pp. 197–217. Springer (2001)
45. Liu, J.S.: Monte Carlo Strategies in Scientific Computing. Springer, New York (2004)
46. Longstaff, F.A., Schwartz, E.S.: Valuing American options by simulation: A simple least-squares approach. *The Review of Financial Studies* **14**(1), 113–147 (2001)
47. Ludkovski, M.: A simulation approach to optimal stopping under partial information. *Stochastic Processes and Applications* **119**(12), 2071–2087 (2009)
48. Mastroeni, L.: Dynamic programming methods for the American option pricing problem with stochastic volatility. *Advances in Mathematical Sciences and Applications* **8**(2), 943–948 (1998)
49. Melenberg, B., Werker, B.J.: A convenient way to characterize equivalent martingale measures in incomplete markets. *Statistical Inference for Stochastic Processes* **2**, 11–30 (1999)
50. Musiela, M., Rutkowski, M.: Martingale Methods in Financial Modeling. Springer, New York (1998)
51. Myneni, R.: The pricing of the American option. *The Annals of Applied Probability* **2**(1), 1–23 (1992)
52. Pan, J.: The jump-risk premia implicit in options: Evidence from an integrated time-series study. *Journal of Financial Economics* **63**, 3–50 (2002)
53. Pastorello, S., Renault, E., Touzi, N.: Statistical inference for random variance option pricing. *Journal of Business & Economic Statistics* **18**(3), 358–367 (2000)
54. Pham, H., Runggaldier, W., Sellami, A.: Approximation by quantization of the filter process and applications to optimal stopping under partial observation. *Monte Carlo Methods and Applications* **11**(1), 57–81 (2005)
55. Pitt, M., Shephard, N.: Filtering via simulation: Auxiliary particle filters. *Journal of the American Statistical Association* **94**(446), 590–599 (1999)
56. Pitt, M., Shephard, N.: Auxiliary variable based particle filters. In: Sequential Monte Carlo Methods in Practice, pp. 177–196. Springer (2001)

57. Rambharat, B.R.: Valuation methods for American derivatives in a stochastic volatility framework. Ph.D. thesis, Carnegie Mellon University, Pittsburgh (2005)
58. Rambharat, B.R., Brockwell, A.E.: Sequential Monte Carlo pricing of American-style options under stochastic volatility models. *The Annals of Applied Statistics* **4**(1), 222–265 (2010)
59. Rambharat, B.R., Brockwell, A.E.: Supplement to “Sequential Monte Carlo pricing of American-style options under stochastic volatility models”. *The Annals of Applied Statistics* (2010)
60. Renault, E., Touzi, N.: Option hedging and implied volatilities in a stochastic volatility model. *Mathematical Finance* **6**(3), 279–302 (1996)
61. Robert, C.P., Casella, G.: Monte Carlo Statistical Methods, first edn. Springer-Verlag, New York (2000)
62. Rogers, L.: Monte Carlo valuation of American options. *Mathematical Finance* **12**(3), 271–286 (2002)
63. Rubenthaler, S., Rydén, T., Wiktorsson, M.: Fast simulated annealing in \mathbb{R}^d with an application to maximum likelihood estimation in state-space models. *Stochastic Processes and their Applications* **119**(6) (2009)
64. Sarwar, G.: An empirical investigation of the premium for volatility risk in currency options for the British pound. *Applied Financial Economics* **12**, 913–921 (2002)
65. Schwartz, E.: The stochastic behavior of commodity prices: Implications for valuation and hedging. *The Journal of Finance* **LII**(3), 923–973 (1997)
66. Seber, G., Wild, C.: Nonlinear Regression, first edn. John Wiley & Sons, Inc., New York (2003)
67. Stein, E., Stein, C.: Stock price distributions with stochastic volatility: An analytic approach. *The Review of Financial Studies* **4**(4), 727–752 (1991)
68. Sullivan, M.A.: Valuing American put options using Gaussian quadrature. *Review of Financial Studies* **13**(1), 75–94 (2000)
69. Touzi, N.: American options exercise boundary when the volatility changes randomly. *Applied Mathematics and Optimization* **39**, 411–422 (1999)
70. Tsitsiklis, J.N., Van Roy, B.: Regression Methods for Pricing Complex American-Style Derivatives. *IEEE Transactions on Neural Networks* **12**(4), 694–703 (2001)
71. Tzavalis, E., Wang, S.: Pricing American options under stochastic volatility: A new method using Chebyshev polynomials to approximate the early exercise boundary. Working paper, Queen Mary, University of London, London (2003)
72. Whaley, R.E.: Derivatives on market volatility: Hedging tools long overdue. *The Journal of Derivatives* (**Fall**), 71–84 (1993)
73. Whittle, P.: Sequential decision processes with essential unobservables. *Advances in Applied Probability* **1**, 271–287 (1969)
74. Wu, A.M.K.: Arbitrage-free evaluation of American-style options on assets with stochastic variance characteristics. In: *Proceedings of EFA 2001 Barcelona Meetings*. Univ. Pompeu Fabra, Barcelona (2001)
75. Yu, J.: On leverage in a stochastic volatility model. *Journal of Econometrics* **127**, 165–178 (2005)
76. Zhang, Z., Lim, K.G.: A non-lattice pricing model of American options under stochastic volatility. *The Journal of Futures Markets* **26**(5), 417–448 (2006)



<http://www.springer.com/978-3-642-25745-2>

Numerical Methods in Finance

Bordeaux, June 2010

Carmona, R.; Del Moral, P.; Hu, P.; Oudjane, N. (Eds.)

2012, XVIII, 474 p., Hardcover

ISBN: 978-3-642-25745-2