

REM WORKING PAPER SERIES

**Least squares Monte Carlo methods in stochastic Volterra
rough volatility models**

Henrique Guerreiro, João Guerra

REM Working Paper 0176-2021

May 2021

REM – Research in Economics and Mathematics

Rua Miguel Lúpi 20,
1249-078 Lisboa,
Portugal

ISSN 2184-108X

Any opinions expressed are those of the authors and not those of REM. Short, up to two paragraphs can be cited provided that full credit is given to the authors.





REM – Research in Economics and Mathematics

Rua Miguel Lupi, 20
1249-078 LISBOA
Portugal

Telephone: +351 - 213 925 912

E-mail: rem@iseg.ulisboa.pt

<https://rem.rc.iseg.ulisboa.pt/>



<https://twitter.com/ResearchRem>

<https://www.linkedin.com/company/researchrem/>

<https://www.facebook.com/researchrem/>

Least squares Monte Carlo methods in stochastic Volterra rough volatility models

Henrique Guerreiro *
hguerreiro@iseg.ulisboa.pt

João Guerra †
jguerra@iseg.ulisboa.pt

ISEG - School of Economics and Management, Universidade de Lisboa

REM - Research in Economics and Mathematics, CEMAPRE
Rua do Quelhas 6, 1200-781 Lisboa, Portugal

May 10, 2021

Abstract

In stochastic Volterra rough volatility models, the volatility follows a truncated Brownian semi-stationary process with stochastic vol-of-vol. Recently, efficient VIX pricing Monte Carlo methods have been proposed for the case where the vol-of-vol is Markovian and independent of the volatility. Following recent empirical data, we discuss the VIX option pricing problem for a generalized framework of these models, where the vol-of-vol may depend on the volatility and/or not be Markovian. In such a setting, the aforementioned Monte Carlo methods are not valid. Moreover, the classical least squares Monte Carlo faces exponentially increasing complexity with the number of grid time steps, whilst the nested Monte Carlo method requires a prohibitive number of simulations. By exploring the infinite dimensional Markovian representation of these models, we devise a scalable least squares Monte Carlo for VIX option pricing. We apply our method firstly under the independence assumption for benchmarks, and then to the generalized framework. We also discuss the rough vol-of-vol setting, where Markovianity of the vol-of-vol is not present. We present simulations and benchmarks to establish the efficiency of our method.

Keywords— VIX, rough volatility, stochastic Volterra models, least squares Monte Carlo, volatility of volatility

*Supported by FCT Grant SFRH/BD/147161/2019.

†Partially supported by the project CEMAPRE/REM-UiDB/05069/2020 - financed by FCT/MCTES through national funds.

1 Introduction

Since the seminal papers [Gatheral et al., 2018] and [Bayer et al., 2016], rough volatility modeling has enjoyed a tremendous focus of interest by the academic community. In such setting, the log-volatility behaves similarly to a fractional Brownian motion with Hurst parameter $H < 1/2$. For this reason, the sample paths are rougher, in the sense that they exhibit a lower Hölder continuity exponent, than those of standard Brownian motion. The rBergomi model, one of the first rough volatility models, adjusts very well to the SP500 index with a small number of parameters. Some of the relevant literature regarding rough volatility includes [El Euch et al., 2018], which provides a micro-structural foundation for rough volatility; [Alòs et al., 2007] and [Fukasawa, 2015], which discuss asymptotic results; as well as a treatment of hedging strategies for the rough Heston model, in [Euch and Rosenbaum, 2017].

A major drawback of the rBergomi model is that it produces flat smiles for the VIX index options market. There are multiple ways to tackle this problem (see [Alòs et al., 2018]). As in the Black-Scholes model volatility was made random to prevent a flat smile for the SP500, it may be natural to introduce stochastic vol-of-vol to produce non-flat VIX smiles. The stochastic Volterra models (SVM) of [Horvath et al., 2020] generalize the rBergomi model in a natural way, seeing it as a truncated Brownian semi-stationary process with constant intermittency. The variance process is

$$v_u = A_0(u) \exp\left(2 \int_0^u g_u(s) \sqrt{\Gamma_s} dW_s\right),$$

for a standard Brownian motion (sBm) W , rough kernel g , vol-of-vol Γ and a deterministic function A_0 . When the vol-of-vol Γ is made non-constant, it is possible to produce upward VIX smiles that resemble those observed in the market.

In the framework of [Horvath et al., 2020], Γ is assumed Markovian and independent of W , which lets the authors develop an efficient pricing mechanism for the VIX. There are however good reasons to pursue a more general framework where these assumptions do not hold. First, a correlation between price and vol-of-vol contributes to upward sloping VIX smiles (see [Gatheral et al., 2020]), which would make the models more flexible. Secondly, there is empirical evidence that variance and vol-of-vol are correlated, and that vol-of-vol may even also display the rough property (see [Da Fonseca and Zhang, 2019]). This more general framework could allow for a more flexible solution to the SP500-VIX joint calibration problem (see [Guyon, 2019] and [Gatheral et al., 2020]).

The main challenge to the more general framework is the computational cost of VIX option pricing. Although Monte Carlo pricing under rough volatility has become more efficient (see [Bennedsen et al., 2017] and [McCrickerd and Pakkanen, 2018]), it is still computationally challenging. Moreover, the proposed Monte Carlo algorithms are very model dependent. To the best of our knowledge, no treatment of the general framework of the SVM has been performed, and no VIX pricing methods have been proposed.

In the absence of a model specific method to price VIX options, we may turn to more model-free methods. The nested Monte Carlo (NMC) method is always available but suffers from severe performance issues due to the extremely high number of simulations required. The classical least squares Monte Carlo (LSMC) method (see [Longstaff and Schwartz, 2001]) can be used to price VIX derivatives for a wide range of models (see [Guo and Loeper, 2018]) but unfortunately works poorly

in a non-Markovian setting, as the state variables will be infinite dimensional. We will explore how the LSMC can be tailored to the SVM to work in this non-Markovian setting. Recently, in [Lelong, 2020], while following a very different approach from ours, the authors propose a variant of the LSMC for Bermudan option pricing that works under non-Markovian conditions. They achieve this by replacing the classical regression by a Wiener chaos expansion.

In this paper, by exploring the structure of the infinite dimensional state variable, we develop a LSMC method for pricing VIX options for the general framework of the SVM. We fully detail the algorithm and present numerical evidence to establish its performance, using the standard NMC as benchmark. We conclude our LSMC is able to accurately price VIX options, achieving the same degree of precision as the NMC at a fraction of the computational cost.

The paper is organized as follows. In Section 2, we will introduce the VIX pricing problem and explain how the NMC and LSMC methods can be used to tackle it. We will draw attention to the challenges to these methods in a non-Markovian setting. In Section 3, we shall then formulate the SVM and discuss its main properties. Namely, we will discuss how to surpass the main obstacles to efficient VIX pricing in the more restricted framework. Afterwards, we will present our LSMC pricing method and show how it can be applied in a more general framework that allows for dependence of vol-of-vol and volatility. Then, in Section 4, we briefly explore the rough vol-of-vol model and discuss the main challenges of the LSMC in that setting. In Section 5, we explore some numerical simulations and compare the results with benchmarks in order to access the performance of the proposed LSMC method. Finally, in Section 6, we present some concluding remarks and suggest further related research problems.

2 VIX Pricing

2.1 The VIX index

Let v denote the variance and $(\mathcal{F}_t)_{t \geq 0}$ be the filtration generated by the processes which drive v . To simplify, we will work under the assumption of zero interest rates. The CBOE VIX index (see [Chicago Board Options Exchange, 2019]), which reflects the market's expectations on future (30 day) volatility, can be defined by the expression

$$VIX_T = \sqrt{\frac{1}{\Delta} \mathbb{E} \left[\int_T^{T+\Delta} v_u du \mid \mathcal{F}_T \right]}, \quad (2.1)$$

where $T > 0$ is a fixed time and $\Delta = 30$ days is the VIX horizon. Let

$$\xi_t(u) = \mathbb{E}[v_u \mid \mathcal{F}_t] \quad (2.2)$$

be the forward variance curve. By a simple application of Fubini's theorem, we can write the VIX in terms of the forward variance curve

$$VIX_T = \sqrt{\frac{1}{\Delta} \int_T^{T+\Delta} \xi_T(u) du}. \quad (2.3)$$

Options and futures can be traded on the VIX, making it a key risk management instrument. If we are to model volatility, we must be able to price VIX options and futures with a reasonable computational cost.

2.2 Nested Monte Carlo

Since the VIX is essentially the square-root of a conditional expectation, there is no direct way to price its options or futures without evaluating the conditional expectation itself. The most general method available for this purpose is the NMC.

In general, let us denote by θ the parameters of a model (e.g. H in the rBergomi model) and by β the initial state variables (e.g. the initial variance in the rBergomi model). Let us assume we possess a procedure *Simulate* that lets us simulate the model's state variables up a time horizon T , given the parameters and the initial state variables. Let us also make the obvious assumption that the variance between T and $T + \Delta$, and hence the integral $\int_T^{T+\Delta} v_u du$, can be obtained from the state variables between times T and $T + \Delta$. Our goal is to create a simulation of N values of

$$Y = \mathbb{E} \left[\int_T^{T+\Delta} v_u du \mid \mathcal{F}_T \right] \quad (2.4)$$

since the VIX can be trivially obtained from it. We would then have a traditional Monte Carlo simulation of the VIX and be able to compute the payoff of any contingent claim.

Notice that in a non-Markovian setting (as in rough volatility models), we must resort to infinite dimensional state variables. In the rBergomi model, this is simply the forward variance curve. In practice, this problem is partially circumvented by projecting such variables into finite dimensional vector spaces.

Bellow we can find a description of the NMC algorithm.

Algorithm 1 Nested Monte Carlo

| | |
|--|---------------------------------------|
| 1: function SIMULATE(θ, β, t) | ▷ Provided by the model |
| 2: function NESTEDMONTECARLO(θ, β, T, N, M) | |
| 3: for $i = 1$ to N do | ▷ Outer paths |
| 4: $\beta'_i \leftarrow$ SIMULATE(θ, β, T) | ▷ Simulates outer path |
| 5: for $j = 1$ to M do | ▷ Inner paths |
| 6: $\beta_j^{(i)} \leftarrow$ SIMULATE(θ, β'_i, Δ) | ▷ Simulates inner path |
| 7: $X_j^{(i)} \leftarrow$ INNERX($\beta_j^{(i)}$) | ▷ Inner estimate |
| 8: $Y_i \leftarrow$ MEAN($X^{(i)}$) | ▷ Averages inner estimates |
| 9: return β', Y | ▷ State variables and target variable |

2.3 Least squares Monte Carlo

As it can be readily noticed by observing the algorithm, the NMC method is extremely slow. This is of course due to the fact that for each outer path, it needs to perform a further Monte Carlo simulation. As a consequence, this algorithm is often impractical. The LSMC method provides a more efficient alternative by attempting to reduce the number of simulations. The algorithm produces N intermediate (potentially noisy) approximations of the target, which are then fed into a regression model. Intuitively, the regression model averages out the noise across the intermediate simulations and is able to learn the true function. This regression model is then used to produce a good

estimate of the target with a lower cost. Depending on the variant, the regression model may or may not use the same data for both fitting and predicting. Moreover, when K outer paths are generated but only N are used for fitting, in the last step of the algorithm we may either use all the K outer paths generated for prediction, or only the $K - N$ that were not used to fit the regression.

Algorithm 2 Least squares Monte Carlo

```

1: function SIMULATE( $\theta, \beta, t$ )                                ▷ Provided by the model
2: function REGRESSIONMODEL( $X, Y$ )                            ▷ Fits a certain regression model
3: function LEASTSQUARESMONTECARLO( $\theta, \beta, T, N, M, K$ )
4:   for  $i = 1$  to  $K$  do                                       ▷ Generate  $K$  outer paths
5:      $\beta''_i \leftarrow$  SIMULATE( $\theta, \beta, T$ )
6:      $\{i_1, i_2, \dots, i_N\} \leftarrow$  SAMPLE( $K, N$ )           ▷ Sample  $N \leq K$  paths
7:     for  $k = 1$  to  $N$  do                                       ▷ Outer paths
8:       for  $j = 1$  to  $M$  do                                       ▷ Inner paths
9:          $\beta_j^{(i_k)} \leftarrow$  SIMULATE( $\theta, \beta'_{i_k}, \Delta$ )  ▷ Simulates inner path
10:         $X_j^{(i_k)} \leftarrow$  INNERX( $\beta_j^{(i_k)}$ )             ▷ Inner estimate
11:         $Y'_{i_k} \leftarrow$  MEAN( $X^{(i_k)}$ )                   ▷ Averages inner estimates
12:     $R \leftarrow$  REGRESSIONMODEL( $\beta', Y'$ )
13:     $Y \leftarrow R(\beta'')$                                        ▷ Regression prediction
14:  return  $Y$ 

```

Increasing LSMC performance

Note that only the N paths used to fit the regression require inner simulations. Thus, one of the main advantages of the LSMC is that there is great flexibility in managing the computational budget, which can be seen as the total number of simulations $K + N \times M$. To start, one might simply choose $K = N$ and see the LSMC as a regularization procedure on top of the NMC, in the sense that a lower number of inner simulations is needed to arrive at the same accuracy. Nevertheless, one of the great advantages of the LSMC is that once the regression model is fitted, it can be used to make fast predictions for unseen values of the predictors. This is especially relevant in our context of option pricing, since out-of-the-money options may require a considerable number of paths K to be evaluated. The cost of such evaluation would be prohibitive using a NMC, where $N = K$. However, provided the regression model is of good quality, it can be used to efficiently generate K paths of the target variable, even if only $N \ll K$ paths were used to learn the mapping. In fact, we can even use the information about the distribution of the predictor contained in the K outer paths to inform the choice of the N paths used for regression fitting.

The choice of N, K and M depends on the nature of the problem and the choice of regression model. For instance, if the regression model is very sensitive to overfitting, we may need to choose a higher value of M and even of N . The complexity of the regression model also plays a role: a simpler regression model may not need to see as many values of the target function, and thus may require a smaller N , allowing for higher M . Note that if the functional form of the target function is known, the LSMC can potentially learn it very efficiently since only a small N might be needed. Finally,

as the dimension of the predictor gets higher, so will N , especially for more complex regression models.

2.4 Challenges of non-Markovian modeling

The regression model presented above is used to approximate a certain random variable Y , using the state variables $\beta \in \mathbb{R}^m$ as predictors. Since we are trying to approximate an arbitrary function f , a classical linear regression model will be inadequate if the function is highly non-linear. A natural strategy is to approximate the projection of f into a finite dimensional vector space. As the Hermite polynomials are an orthogonal basis of L^2 , we can write

$$f = \tilde{f} + \varepsilon = \sum_{k=1}^n \alpha_k H_k + \varepsilon, \quad (2.5)$$

where $(H_k)_{k \leq n}$, are the n Hermite polynomials with degree no greater than d in \mathbb{R}^m , and ε is an error term. We may find α_k using classical linear regression, with predictors $H_k(\beta_i)_{k=1}^n, i = 1, \dots, N$. The price we have to pay, however, is the exponentially increasing number of predictors with respect to the number of state variables m .

In Markovian models, it is easy to encode all information known at time T in a vector with only a few entries. In non-Markovian models, however, because we have to project the infinite dimensional state variable of the process into a finite dimensional vector space, we will have a vector β_i with tens, or even hundreds, of entries. This challenge turns the ‘‘classical’’ LSMC as described above, unfeasible.

In the next section, we will explore a very general family of models, of which the rBergomi model is a special case. The LSMC can be adapted to become a valid alternative for VIX option pricing in these models when some of the ‘‘nicer’’ analytical conditions are removed.

3 Stochastic Volterra Models

The stochastic Volterra models of [Horvath et al., 2020] provide a promising framework to tackle the SP500/VIX Joint Calibration Problem. In these models, the variance behaves as a truncated Brownian semi-stationary process (TBSS). They generalize the rBergomi model and are able to produce an upward slopping VIX smile. We start with the definition of the truncated Brownian semi-stationary process (TBSS), following [Barndorff-Nielsen and Schmiegel, 2009] and [Bennedsen et al., 2017].

Definition 3.1

We say X is a truncated Brownian semi-stationary process if

$$X_u = \int_0^u \sqrt{\Gamma(s)} g(u-s) dW_s, \quad (3.1)$$

for a sBm W , a predictable locally bounded process Γ , and a Borel measurable function g . The processes Γ and W may possibly be dependent.

In order to simulate paths of a TBSS, we may use the Hybrid scheme, introduced in [Bennedsen et al., 2017], which is able to deal with the singularity posed by rough

3 STOCHASTIC VOLTERRA MODELS

kernels. The main idea of the method is to approximate the integral near the singularity by that of a power function, while using regular Riemann sums away from the singularity. The final approximation then involves Wiener integrals of a power function as well as a Riemann sum, and hence the name *hybrid* scheme. In a SVM, the log-variance is modeled by a TBSS:

Definition 3.2 (Stochastic Volterra model)

Let B and W be two ρ -correlated Brownian Motions, with $-1 < \rho < 1$. We say that (S, v) follows a stochastic Volterra model if

$$dS_t = S_t \sqrt{v_t} dB_t, \quad (3.2)$$

where the variance process is given by

$$v_u = A_0(u) e^{2X_u} \quad (3.3)$$

and X_u is a TBSS given by

$$X_u = \int_0^u \sqrt{\Gamma_s} g_u(s) ds. \quad (3.4)$$

The rough kernel g is in our setting given by

$$g_u(s) = (u - s)^{H-1/2} \quad (3.5)$$

for Hurst parameter H , with $0 < H < 1/2$ and a deterministic function A_0 . The vol-of-vol process Γ starts at a given constant $\gamma > 0$.

Remark 1. In the paper [Horvath et al., 2020], the vol-of-vol process is assumed to be independent from W and Markovian. In the definition above we are intentionally dropping those assumptions. This leads to a more general framework, with the obvious cost of analytical and computational difficulties. Attempting to overcome those difficulties is the main subject of our paper.

The rBergomi model of [Bayer et al., 2016] can be seen as a particular case of the SVM, one in which the vol-of-vol is constant and given by

$$\Gamma(s) \equiv \gamma = \frac{\eta^2 H}{2}, \quad (3.6)$$

for a fixed $\eta > 0$.

The rBergomi model is able to calibrate the SP500 very well but cannot calibrate the VIX smile since the market exhibits upward slopping smiles whilst the rBergomi model, by virtue of its constant vol-of-vol, produces approximately flat VIX smiles. The extra randomness introduced by stochastic vol-of-vol provides the flexibility needed for an upward slopping VIX smile, as it was showed in [Horvath et al., 2020].

In order to develop a successful VIX option pricing method in the general framework, let us try to find an infinite dimensional representation for the SVM.

Proposition 3.1

For $u \geq t > 0$, let

$$A_t(u) = \frac{\xi_t(u)}{h_t(u)}, \quad (3.7)$$

where

$$h_t(u) = \mathbb{E}[E_{t,u}(u) | \mathcal{F}_t] \quad (3.8)$$

and

$$E_{p,q}(u) = \exp\left(2 \int_p^q g_u(s) \sqrt{\Gamma_s} dW_s\right). \quad (3.9)$$

We have the time-invariant decomposition

$$v_u = A_0(u)E_{0,u}(u) = A_t(u)E_{t,u}(u) \forall u \geq t > 0.$$

In addition, if Γ is Markovian, the process A_t provides an infinite dimensional Markovian representation of the SVM in the sense that conditional on A_t and Γ_t , v_u is independent of \mathcal{F}_t .

The (simple) proof of this fact can be found in Appendix A.

Remark 2. In the rBergomi model, $\Gamma \equiv \gamma$, constant. In such setting

$$\begin{aligned} h_t(u) &= \mathbb{E}\left[\exp\left(2\sqrt{\gamma} \int_t^u g_u(s) dW_s\right) \mid \mathcal{F}_t\right] \\ &= \exp\left(2\gamma \int_t^u (u-s)^{2H-1}\right) \\ &= \exp\left(\frac{1}{2}\eta^2(u-t)^{2H}\right), \end{aligned}$$

where we have used Eq. (3.6). This means $A_t(u)$ and $\xi_t(u)$ are equivalent in the rBergomi model, whereas in the SVM we need to compute the (random) conditional expectation $h_t(u)$ in order to relate them.

Remark 3. Note that regardless of whether Γ is Markovian, the ‘‘hard’’ part of VIX pricing is the curve h_T . Indeed,

$$\Delta VIX_T^2 = \int_T^{T+\Delta} \xi_T(u) du = \int_T^{T+\Delta} \frac{\xi_0(u)}{h_0(u)} E_{0,T}(u) h_T(u) du.$$

The value $E_{0,T}(u)$ acts as a state variable and can be computed simply using Riemann sums. The initial forward variance curve ξ_0 can be observed from the market or assumed flat and taken as a parameter. As for h_0 , it can be written as follows:

$$h_0(u) = \mathbb{E}[E_{0,u}(u)] = \mathbb{E}[E_{0,T}(u) \mathbb{E}[E_{T,u}(u) \mid \mathcal{F}_T]] = \mathbb{E}[E_{0,T}(u) h_T(u)].$$

Thus, since we know $E_{0,T}(u)$ and $h_T(u)$, we may estimate $h_0(u)$ by averaging the product of these two variables. Alternatively, h_0 can be estimated through a (simple) Monte Carlo simulation.

3.1 Independent and Markovian vol-of-vol

In the discussion of the Markovian and independent vol-of-vol process, we follow [Horvath et al., 2020]. Let us assume Γ is Markovian and independent of the sBm W . Then, for any $u \geq t > 0$,

$$h_t(u) = \mathbb{E}[E_{t,u}(u) \mid \Gamma_t]. \quad (3.10)$$

By conditioning on the full path of Γ between t and u , $E_{t,u}(u)$ is conditionally log-normal with mean zero, so that

$$h_t(u) = \mathbb{E}\left[\exp\left(2 \int_t^u g_u^2(s) \Gamma(s) ds\right) \mid \Gamma_t\right]. \quad (3.11)$$

If Γ is assumed to be exponentially affine, along some integrability conditions, the above conditional expectation can be written in terms of the solution to a certain ODE and the values Γ_0 and Γ_T . For details, we refer to Appendix A. This leads to an efficient framework for VIX pricing since there is no need for inner simulations in order to compute the conditional expectation. An example of a process that satisfies the above conditions, as provided in [Horvath et al., 2020], is the CIR process:

$$d\Gamma_t = \theta(m - \Gamma_t) dt + \delta\sqrt{\Gamma}dZ_t, \quad (3.12)$$

where $\theta, m, \delta > 0$ and Z is a sBm independent of W .

3.2 Dropping independence

There is some evidence that dropping one (or both) the assumptions of independence or Markovianity is justified. First, the vol-of-vol process might be itself a rough volatility process, which suggests a non-Markovian vol-of-vol (see [Da Fonseca and Zhang, 2019]). Additionally, vol-of-vol is higher in periods of crisis, suggesting a negative correlation with the asset price, and this contributes to upward slopping VIX smiles (see [Gatheral et al., 2020] and [Goutte et al., 2017]). Even if we only wished to correlate vol-of-vol and price, this might require variance and vol-of-vol to be correlated. In order to see this, let us take the CIR example above and assume that B, W and Z , which are the standard Brownian motions driving the price, volatility and vol-of-vol, respectively, form a 3-dimensional Brownian motion with correlation matrix

$$\Sigma = \begin{bmatrix} 1 & \rho & \rho_S \\ \rho & 1 & \rho_V \\ \rho_S & \rho_V & 1 \end{bmatrix}.$$

The determinant of this matrix is of course

$$1 - \rho^2 - \rho_S^2 - \rho_V^2 + 2\rho\rho_S\rho_V.$$

If variance and vol-of-vol are independent, $\rho_V = 0$ and the determinant becomes

$$1 - \rho^2 - \rho_S^2.$$

But since we know from empirical data that $\rho \approx -1$, for the matrix to be positive definite we must have $\rho_S \approx 0$.

The question now is: can we generalize the above framework by dropping one of the two crucial assumptions (independence and Markovianity) whilst keeping the VIX pricing methods with a reasonable running time?

3.3 Least squares Monte Carlo

Let us drop the assumption that Γ and W are independent but assume Γ is still Markovian. Let us take the infinite dimensional state variable h_T . Then, we fix a time grid $(u_j)_{j=1}^n$ between T and $T + \Delta$ and obtain the projection $Y = (h_T(u_j))_{j=1}^n$. We will use the LSMC method to approximate Y . Since Γ is Markovian, we have

$$Y_j = h_T(u_j) = \mathbb{E} \left[2 \int_T^{u_j} \sqrt{\Gamma_s} g_{u_j}(s) \mid \Gamma_T \right] = f_j(\Gamma_T), \quad (3.13)$$

for some deterministic function f_j . Thus, we may write this as a classical multivariate regression problem:

$$Y = f(\Gamma_T) = \tilde{f}(\Gamma_T) + \varepsilon, \quad (3.14)$$

where \tilde{f} is the deterministic approximation function and ε is the stochastic error term. Note that the number of predictors is simply one (the vol-of-vol process at time T). Increasing the number of grid points will only increase the size of the output vector, leading to a linear, and not exponential, increase in complexity.

In order to take maximum advantage of the LSMC, especially in this one-dimensional problem, we allow for $N < K$ (see Section 2.3). In order to use the information contained in the K outer paths, we sample N out of the K generated in order to fit the regression model.

Fix a time grid $(u_j)_{j=1}^n$ between T and $T + \Delta$. The LSMC method for SVM is then as follows.

1. Run an outer simulation and obtain K realizations of the state variables Γ_T and $[E_{0,T}(u_j)]_{j=1}^n$.
2. Of these K outer paths, take a sample of $N \leq K$.
3. For each outer value $\Gamma_T^{(i)}$ of vol-of-vol, run an inner simulation to obtain $[E_{T,u_j}^{(i)}(u_j)]_{j=1}^n$, which is a TBSS with initial vol-of-vol $\Gamma_T^{(i)}$.
4. Compute the mean over inner simulations and obtain intermediate estimates y'_j of $h_T(u_j)$ for $j = 1, \dots, n$.
5. Fit the regression model with predictor Γ_T and dependent vector $(y'_j)_{j=1}^n$.
6. Use the fitted regression model to obtain the final estimates of $[h_T(u_j)]_{j=1}^n$ by evaluating the K initially generated outer paths.
7. Obtain an estimate for $h_0(u_j)$ by averaging the product $E_{0,T}(u_j)h_T(u_j)$, for each $j = 1, \dots, n$.
8. Compute the forward variance curve using the formula

$$\xi_T(u_j) = \frac{\xi_0(u_j)}{h_0(u_j)} E_{0,T}(u_j) h_T(u_j), j = 1, \dots, n. \quad (3.15)$$

9. Compute the VIX value by approximating the integral in Eq. (2.3) using a classical quadrature method.

Remark 4. We note that the choice of regression model need not be the linear regression on a function basis like the Hermite polynomials. Indeed, non-linear non-parametric methods such as Random Forests and Neural Networks can be used. For more details about Random Forests, we refer to Appendix B.

4 Rough vol-of-vol

In [Da Fonseca and Zhang, 2019], the authors have provided empirical evidence that volatility of volatility is also rough. They took high frequency data for the VIX index and computed its volatility, which acts as volatility of volatility. Then they showed its logarithm resembles a fractional Brownian motion with Hurst parameter less than $1/2$, similarly to what was done in [Gatheral et al., 2018] to establish roughness of volatility.

Moreover, they estimated a positive correlation between volatility and volatility of volatility.

The (extended) SVM framework can accommodate for a rough vol-of-vol. Indeed, we may simply consider the vol-of-vol Γ as a rough log-TBSS with constant volatility

$$\Gamma_s = Q_0(s)I_{0,s}(s), \quad (4.1)$$

$$I_{0,s}(s) := \exp\left(2\sqrt{\nu} \int_0^s g_s(t) dZ_t\right), \quad (4.2)$$

where Z is a sBm correlated with W and B , $\nu > 0$ is a constant and Q_0 is a deterministic function. Since Γ is a log-TBSS with constant volatility, using Proposition 3.1 and proceeding as in Remark 2, we know the law of Γ only depends on \mathcal{F}_T through

$$\zeta_T(u) := \mathbb{E}[\Gamma_u | \mathcal{F}_T]. \quad (4.3)$$

Let us then assume that the initial curve is flat $\zeta_0(u) \equiv \zeta_0$, where ζ_0 is treated as a model parameter to be chosen. Then, as in the rBergomi model, we can obtain ζ_T by the formula

$$\zeta_T(u) = q_T(u)\zeta_0(u)I_{0,T}(u), \quad (4.4)$$

where q_T is an analytically known deterministic function. In this model, the vol-of-vol Γ is neither Markovian nor independent of W . It therefore poses a formidable computational task. As seen above, in order to compute a VIX path we need the curve

$$h_T(u) = \mathbb{E}\left[\exp\left(2\int_T^u g_u(s)\sqrt{\Gamma_s} dW_s\right) | \mathcal{F}_T\right] \quad (4.5)$$

$$= F(\zeta_T(s)_{T \leq s \leq u}), \quad (4.6)$$

since Γ only depends on \mathcal{F}_T through ζ_T .

If we wish to run a simulation of this model and the relevant state variables we may proceed as follows:

1. Produce paths for the three dimensional Brownian Motion (B, W, Z) , where the covariance matrix is given.
2. Using the paths of Z , apply the hybrid scheme method (note constant volatility of the TBSS) to obtain paths for Γ and $I_{0,T}(t)$ for $T \leq t \leq u$ with t in some time grid.
3. Apply the updating formula (4.4) and obtain

$$\zeta_T(t) = \zeta_0(t)q_T(t)I_{0,T}(t) = \zeta_0q_T(t)I_{0,T}(t).$$

4. For each obtained path of Γ , apply the hybrid scheme to compute the integral $\int_0^u g_u(s)\sqrt{\Gamma_s} dW_s$ and then proceed as usual.

The LSMC method for rough vol-of-vol could then be applied as follows:

1. Run an outer simulation and obtain K realizations of the state variables $\zeta_T(u_j)$ and $E_{0,T}(u_j)$ for $j = 1, \dots, n$.
2. For each i of $N \leq K$ selected outer paths, apply the procedure above to obtain M inner paths of $[E_{T,u_j}^{(i)}(u_j)]_{j=1}^n$, which is a TBSS with rough volatility, where the (projection of the) initial curve is $[\zeta_T^{(i)}(u_j)]_{j=1}^n$. Again, we may choose $N < K$ if needed.

3. Compute the mean over inner simulations and obtain intermediate estimates y'_j of $h_T(u_j)$ for $j = 1, \dots, n$.
4. Fit the multivariate regression model with predictors $[\zeta_T(u_j)]_{j=1}^n$ and dependent vector $(y'_j)_{j=1}^n$.

The remaining steps are as in the Markovian vol-of-vol case. The added difficulty now is that the trick we used to escape an infinite dimensional predictor does not work, since even h_T depends on \mathcal{F}_T through an infinite dimensional object. Thus, the size of our predictor makes the usual linear regression on Hermite polynomials unpractical. In order to price VIX options in this setting one needs to consider the whole path ζ_T as a predictor.

Although a full non-Markovian method for VIX option pricing is outside the scope of this article, we would like to point out some promising approaches. Since the structure of the dependence between h_T and ζ_T is unknown, general high dimensional machine learning methods stand out as candidates for regression models in the LSMC. Of these we may name Random Forests, which are quite flexible and simple to implement, and Neural Networks, which are a focus of research and have been used before in the context of rough volatility (see [Stone, 2019], [Bayer et al., 2019] or [Horvath et al., 2019]). Finally, we would like to point to the Wiener Chaos approach of [Lelong, 2020], which can in principle be extended to rough volatility models.

5 Numerical Experiments

The NMC for VIX pricing has three main sources of error: the number of time steps used for the discretization grid, the number of outer paths and the number of inner paths. In the context of LSMC, an extra source of error – the regression model – has to be considered. For each source, the error can be reduced by paying the corresponding computational price. Our main finding is that the LSMC allows to obtain accurate estimates at a reasonable computing cost, while the NMC requires a much higher computational cost in order to obtain the same accuracy. In fact, in some models, accurate pricing by NMC would require such a high computational cost as to be made impractical.

As seen in Algorithm 2, we use only $N < K$ outer paths in order to fit our regression model but then use K outer paths for the VIX by using the already fitted regression model. This key aspect of the LSMC allows it to vastly outperform the NMC.

Thus, in order to assess the quality of our LSMC, we proceeded as follows. First, we produced $K = 50,000$ (outer) simulations of the model's state variables given the initial conditions and the parameters. Then, for various configurations of N and M , we randomly select $N \leq K$ outer paths and produce M inner simulations for these N outer paths, which we use to fit various LSMC models for different regression model choices. The choice of N and M might depend on the regression model choice and the specification of the dynamics of Γ . For the NMC obviously $N = K$.

In line with LSMC practices, we consider the linear regression of Hermite polynomials as a regression model choice. The maximum degree of Hermite polynomials was set to 3. Also, for their generality and wide applicability in machine learning, we also use Random Forests and Neural Networks. In order to avoid overfitting, the tree maximum depth was set to 5. All the remaining Random Forest hyper-parameters were set to

the defaults of the *RandomForestRegressor* class of the *sklearn.ensemble* Python library. For the Neural Network, we used three hidden layers of size 32, with a batch size of 32 and a learning rate of 10^{-3} . The sample was divided in training in validation, with 80% used for training and 20% for validation. Finally, we incorporate a simple Linear Regression in our analysis since h_T is log-linear in the setting of Section 3.1. This simple toy example will also show us the power of the LSMC when the functional form the target function is known. The parameters of the model were chosen based partly on known best practices, as well as some numerical tests, and for this reason they may be changed if the underlying model dynamics are changed. A further hyper-parameter optimization routine, based per example in cross validation, can be used to find the adequate hyper-parameters if one wishes to further increase the performance of LSMC.

In very general models, especially Random Forests but also Neural Networks, it is harder to generalize to tail values of the input. Thus, when sampling N outer paths out of the initially K generated outer paths, we extract a stratified sample. We split the input data in evenly spaced sub-intervals (strata) and sample from each one until a certain minimum sample goal is achieved or all points of the strata are sampled. The rest of the points are sampled uniformly to preserve the distribution of Γ_T . Thus, we make sure regions with higher probability mass have more weight while at the same time ensuring the regression model can cover the whole sample space.

The time horizon considered is $T = 7$ days, where the VIX horizon is, as usual, $\Delta = 30$ days. The number of time points per day is set to $n_d = 7$, so that it makes only a small contribution to the error. The initial forward variance curve is assumed flat $\xi_0(u) \equiv v_0$ and we set $v_0 = 0.013$. The Hurst index is set to be $H = 0.1$. The correlation between B and W is $\rho = -0.95$ and the initial vol-of-vol is $\gamma = 0.05$. The vol-of-vol follows the CIR model where the parameters are chosen as follows: $\theta = 0.4, \delta = 0.8, \kappa = 0.8125$. We choose these parameters so that they both satisfy the conditions in [Horvath et al., 2020], to allow for benchmarks, while also producing reasonable VIX smiles. The time grid for $[T, T + \Delta]$ is as in the Hybrid Scheme with $\kappa = 1$:

$$u_j = T + \frac{j}{365n_d}.$$

All numerical simulations are performed on a Linux machine (Ubuntu 20.4), with AMD Ryzen 7 3800X CPU with 16 threads. The Python language was used, resorting to *numpy* whenever possible for C-like speed. In order to train regression models, we used the *scikit-learn* and *pytorch* libraries. As most numerical tasks involved are *embarrassingly parallel*, parallel computing was used to reduce running times.

5.1 Markovian and independent vol-of-vol

Let us first consider the case where the vol-of-vol Γ is independent of W and Markovian. In this case, we may use the approach of [Horvath et al., 2020] to know the true value of $h_T | \Gamma_T = \Gamma_T^{(i)}$ for each outer simulation i . This allows us to compute useful metrics for the LSMC method. Moreover, since the independence assumption is not used by the LSMC, we are lead to believe the quality of the algorithm generalizes well to the dependent case.

5.1.1 Exploring the data

In order to start the data analysis, for a fixed u , we provide a scatter plot of the estimates for $h_T(u)$ produced by the NMC and the different variants of the LSMC algorithm in \log scale. The degree of noise is larger for larger u . We choose $u = u_{141}$, about two thirds of the way between T and $T + \Delta$. In Fig. 1, we can compare how the various regression models adjust to the data, while also getting a feel of how noisy the NMC data is with the maximum number of inner paths considered of $M = 1000$. The proxy for the true value of h_T , coming from the ODE method described in Appendix A, is displayed for comparison in a solid line. The NMC, as expected, displays a great amount of noise.

In the independent Markovian case we know the function h_T exactly:

$$h_T(u) = \exp(\psi(u - T)\Gamma_T + \phi(u - t)), \quad (5.1)$$

where ψ and ϕ are as in Theorem A.1. The function is log-linear and thus it is no surprise both the linear model and the Hermite present very good fits. Both the Random Forest and Neural Network present satisfactory fits. We will test these methods with a non-linear function in Section 5.1.3.

5.1.2 Variation of the number of inner paths

We now wish to know the behavior of each algorithm as the number of paths increases. Since the configuration of N and M is different for each LSMC variant and NMC, and different regression models have different computing times, we found best to use running time to compare the performance of the models.

In Fig. 2, we observe the evolution of the root mean squared Error of h_T . As could be inferred from Fig. 1, the NMC presents much larger errors for the same computing cost. In all LSMC variants, the error quickly diminishes as the number of paths increases.

The variance present in the estimates of h_T for the NMC will increase the VIX variance and consequently the price of VIX options. For this reason, we observe a consistent positive bias in the NMC estimates for VIX options and implied volatilities. This is apparent in Fig. 3. Again, even the more general LSMC variants exhibit fast convergence to the true value. The behavior is similar in Fig. 4, where we can observe the evolution of the RMSE with respect to the entire implied volatility smile.

5.1.3 Non-linear data

As noted before, the LSMC Hermite and LSMC Linear methods are privileged because the true function h_T is log-linear in Γ_T . To test how these methods perform when the function to learn is not log-linear, we learn $(\log h_T)^3$ instead of $\log h_T$. In Fig. 5, we can see the behavior of each LSMC in this case. The Random Forest, because of its immense flexibility ends up being sensitive to outliers in sparsely sampled regions. Nevertheless, this does not compromise the learning in more densely sampled regions, which have more importance.

In Fig. 6, we can more accurately access the performance of each LSMC variant in this non-linear setting. The linear regression obviously does not (can not) capture the

5 NUMERICAL EXPERIMENTS

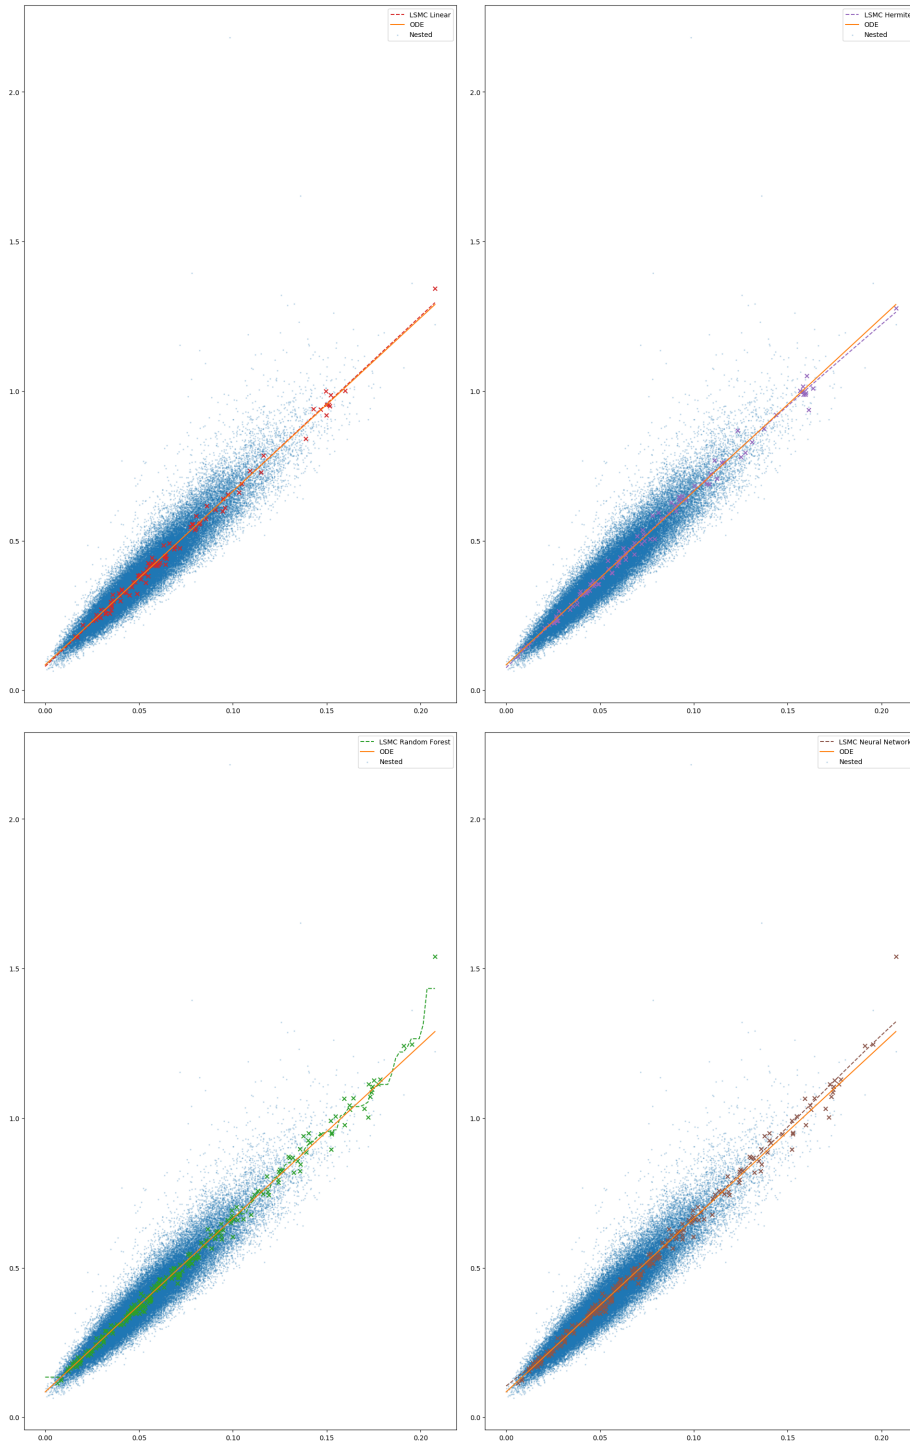


Figure 1: Estimates of h_T for various LSMC variants

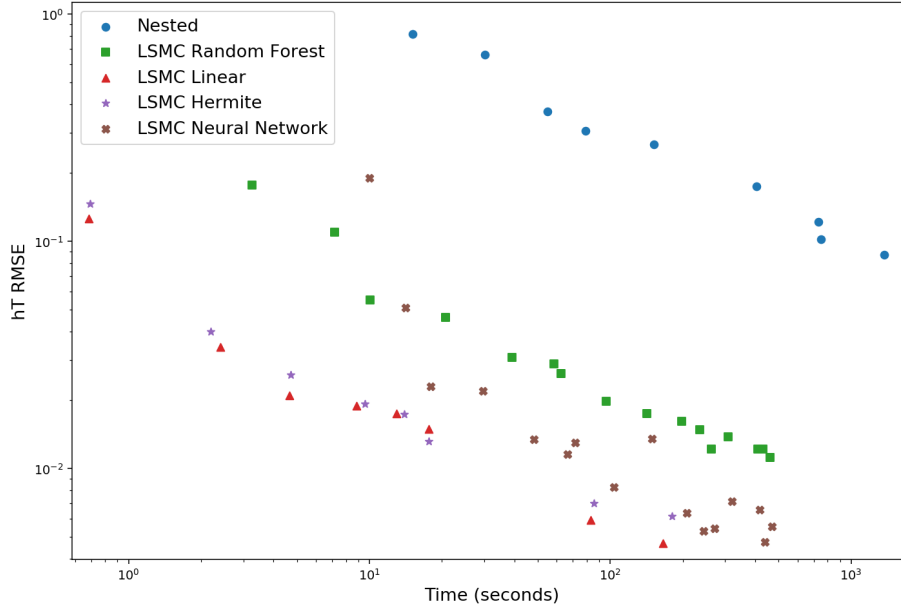


Figure 2: Evolution of the RMSE of h_T with running time

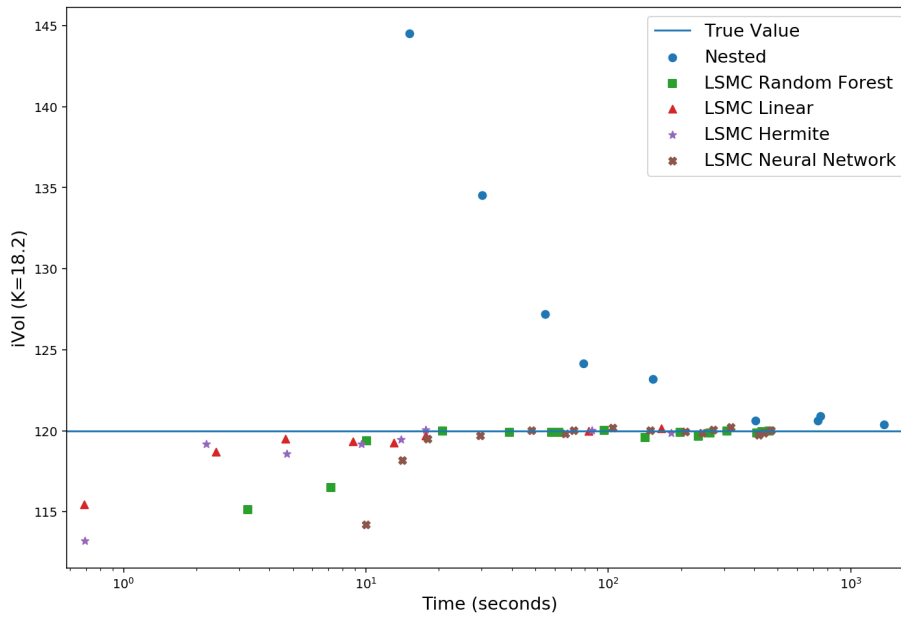


Figure 3: Convergence of implied volatility

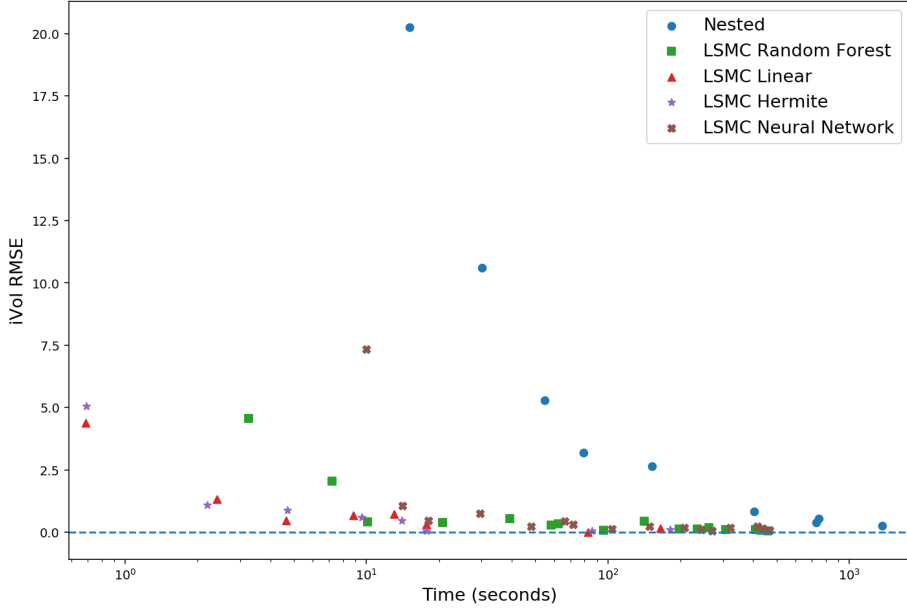


Figure 4: Evolution of the RMSE of implied volatility with running time

behavior of h_T and thus does not converge to the true value. Both the LSMC Random Forest and the LSMC Hermite still provide reasonably accurate estimates, indicating their reliability even when the true structure of the mapping h_T is unknown.

The LSMC Neural Network exhibits some variance in its estimates and a slower convergence. Perhaps it is possible to tailor the Neural Network architecture to the problem at hand and improve its performance. However, optimizing Neural Network architecture is outside the scope and goal of this paper. Because Neural Networks are more complex, hyper-parameter optimization is thus harder when compared to other variants of the LSMC.

5.1.4 Performance metrics

If we wish to compare LSMC methods with the NMC, we can compare their running times. In Fig. 7, we compute the ratio of the NMC running time and each LSMC running time for a given goal of iVol RMSE. Because the LSMC Random Forest and LSMC Neural Network have time consuming regressors and require higher values of N , their performance is not as high as the LSMC Hermite. Nevertheless, they still outperform the NMC by a considerable factor. The Neural Network factor ranges from 3.88 to 28.67, while the Random Forest ranges from 6.93 to 72.57. The LSMC Hermite is the best performing method, preserving adaptability and outperforming the NMC by a huge factor.

Finally, in Fig. 8, we can observe the different volatility smiles produced by the different methods for different number of generated paths. We present the base 10 logarithm of the total number of paths $N \times K$ in the labels.

We may thus conclude the NMC, although valid, is very computationally expensive,

5 NUMERICAL EXPERIMENTS

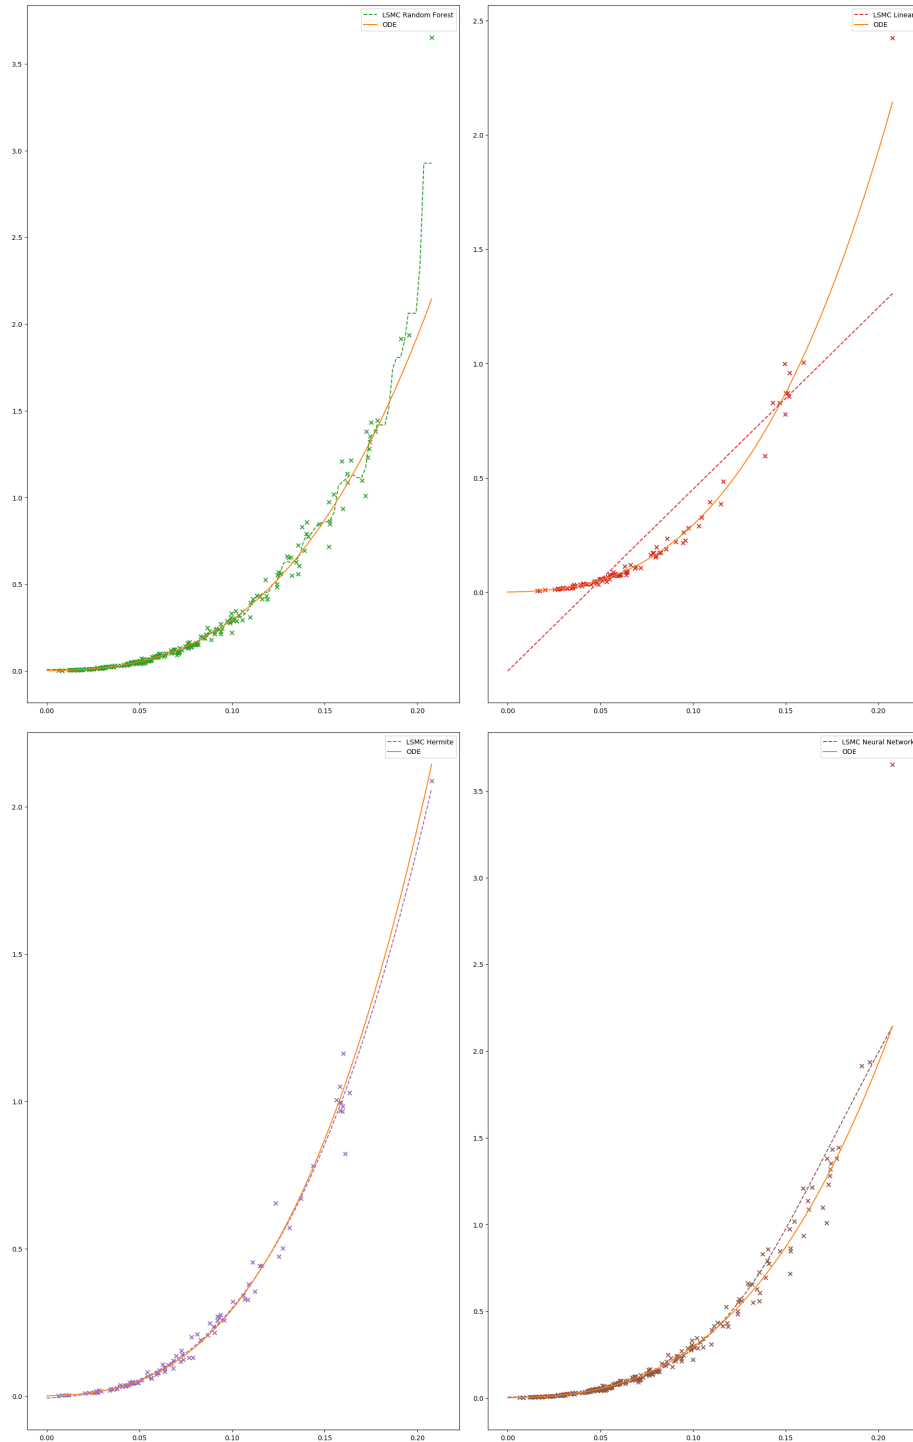


Figure 5: Estimates for h_T for various LSMC variants – non-linear case

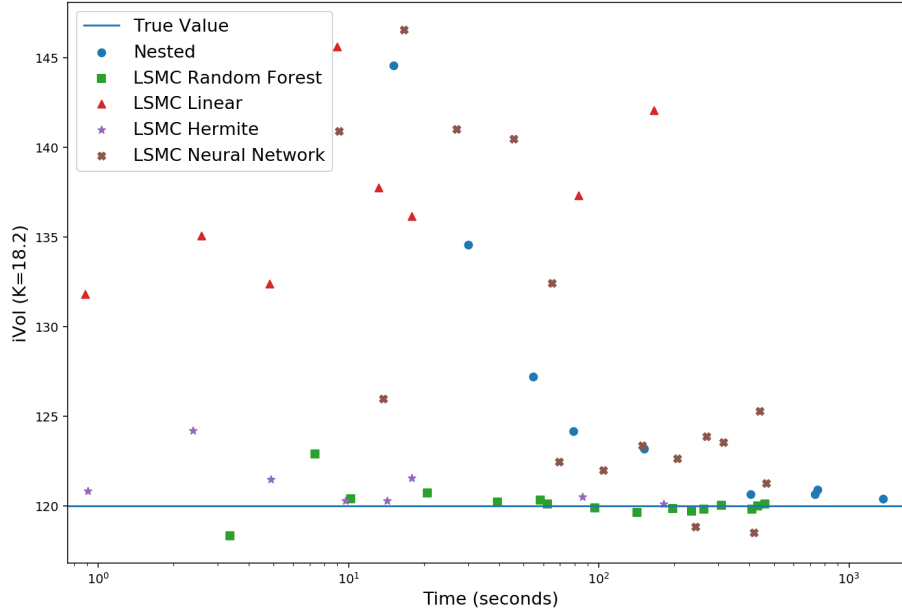


Figure 6: Convergence of implied volatility – non-linear case

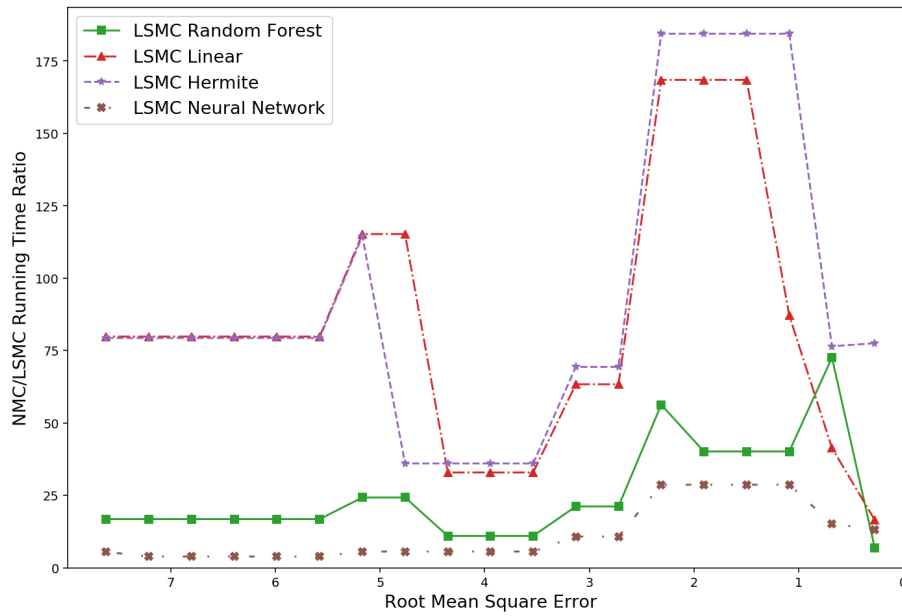


Figure 7: Relative comparison of running times

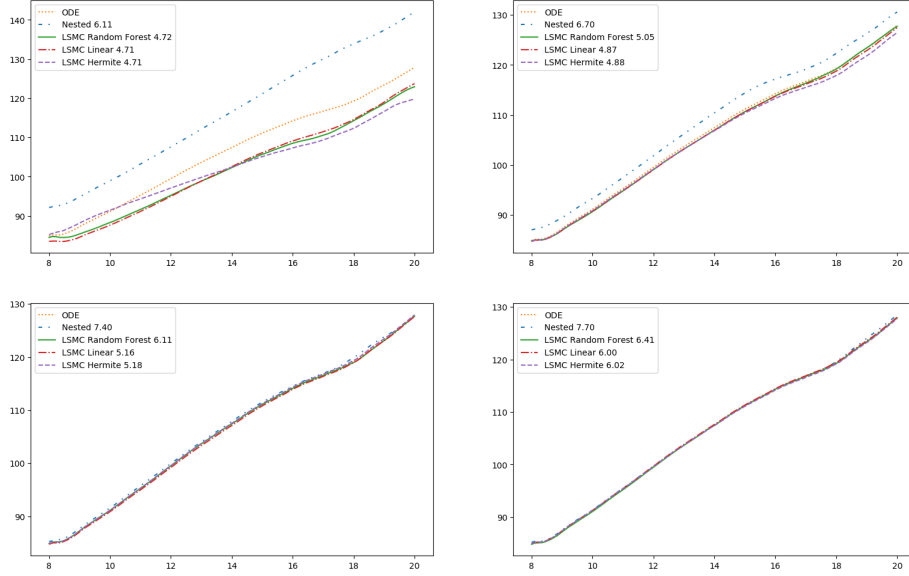


Figure 8: Comparison of produced iVol Smiles

and is deemed impractical in a production environment. On the other hand, the LSMC method, across its different variants, performs well in pricing VIX derivatives, being able to produce accurate results by a fraction of the computational cost of the NMC. In particular, the LSMC Hermite, frequently deployed in a LSMC setting, shows both fast computing time and adequate robustness.

5.2 Dependent vol-of-vol

Let us now consider the case where the vol-of-vol process Γ is no longer independent of W but still Markovian. In this case, the method presented in [Horvath et al., 2020] is no longer available, so we cannot compute metrics such as the RMSE. This is, of course, where the LSMC shines.

Although now $h_T(u)$ is not necessarily log-linear in Γ_T , we will still consider $\log h_T$ as the function to predict and include the simple linear regression model in our set of LSMC methods. This lets us examine how well h_T resembles a log-linear function as in the previous case. If h_T happens to still (even approximately) keep the log-linear property, the simple linear regression will be a useful LSMC method. It is worth noting that if were to drop the Markovianity assumption and $\log h_T$ was a linear functional of the infinite dimensional object that encodes the information at \mathcal{F}_T , the VIX pricing problem via LSMC would be reduced to fitting a high dimensional linear regression. In order to model dependence, we simply let the sBm (B, W, Z) form a 3-dimensional sBm with correlation matrix

$$\Sigma = \begin{bmatrix} 1 & \rho & \rho_S \\ \rho & 1 & \rho_V \\ \rho_S & \rho_V & 1 \end{bmatrix}.$$

We set $\rho_S = -0.9$ and $\rho_V = 0.9$ so that there is strong dependence. The rest of the parameters remain as in the previous section.

5.2.1 Exploring the data

When compared to the previous case, there is somewhat higher variance because of the high correlation, as can be observed in Fig. 9.

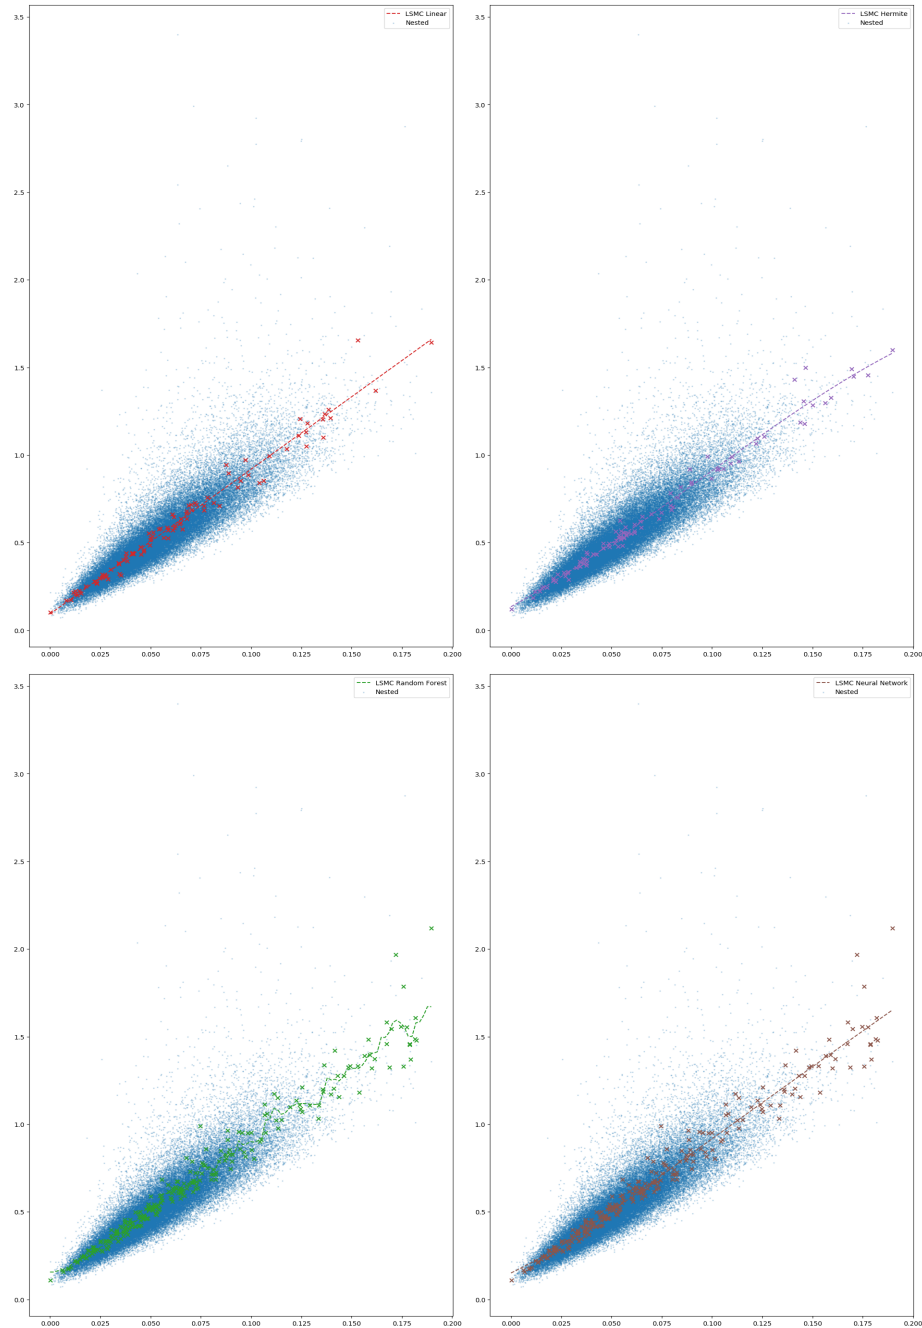


Figure 9: Estimates of h_T for various LSMC variants – dependent Γ

5.2.2 Convergence of implied volatility

We are not in possession of the true value of h_T for a given outer path, but given the results in the previous section, we expect the LSMC methods to converge to the true value. It is not guaranteed the simple linear regression LSMC will converge to the same value, since there is no guarantee that h_T is log-linear in Γ_T , but from the data exploration conducted previously, it looks it might at least be well approximated by a log-linear function.

Because the data exhibits more variance, we expect slower convergence and a higher bias of the NMC when estimating implied volatility. Nevertheless, we expect both the NMC and the LSMC to converge towards the same value. This behavior is confirmed in Fig. 10.

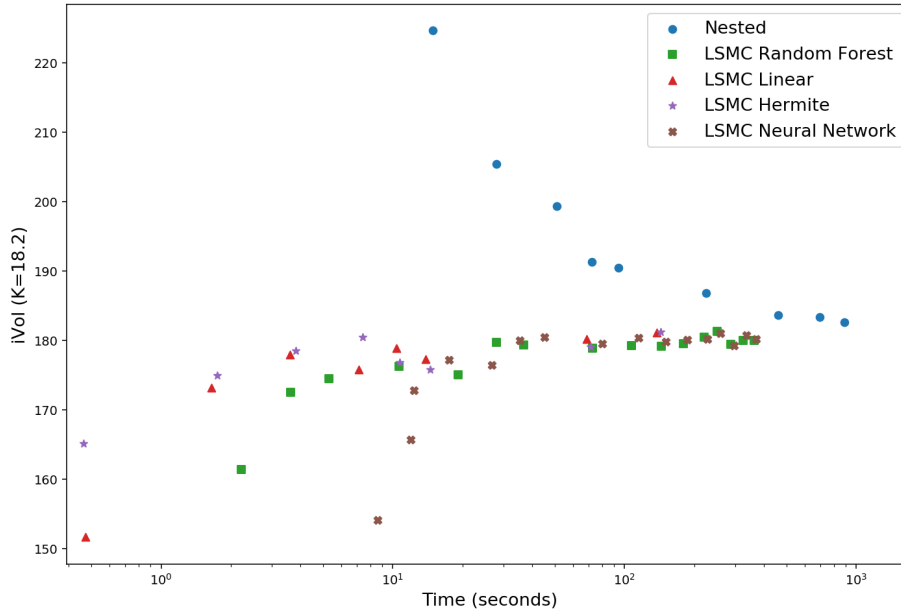


Figure 10: Convergence of implied volatility – dependent Γ

5.2.3 Generated smile

In Fig. 11, we can see the smiles generated in the dependent case, which we may compare to Fig. 8, to see the impact of adding correlation. Although slowly than in Section 3.1, the generated smiles approach a consistent limit as the computational budget increases.

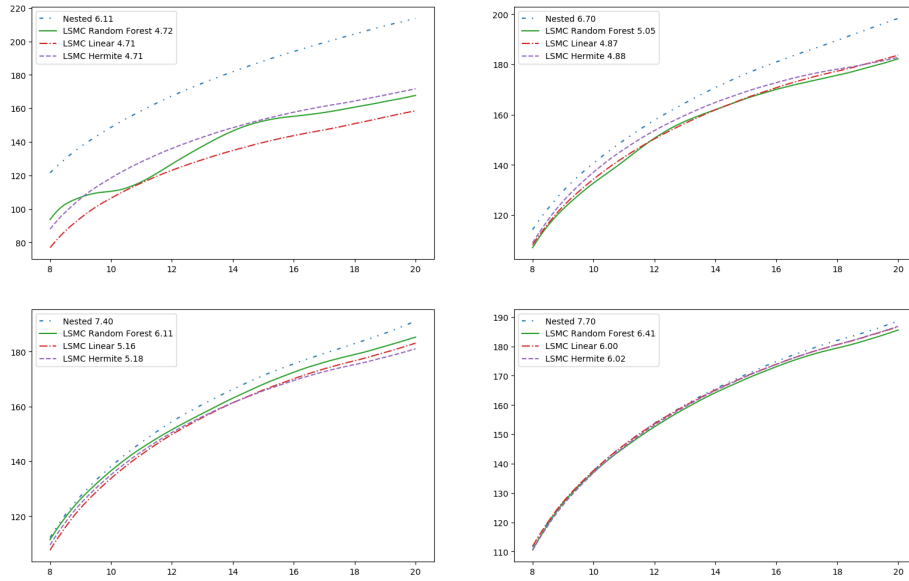


Figure 11: Comparison of produced implied volatility smiles – dependent Γ

6 Conclusion and Further Research

In this paper, we proposed a new method for VIX derivatives pricing in stochastic Volterra models. This method works even in the more generalized framework of SVM, where to the best of our knowledge no VIX pricing method had been proposed. We attested to the good performance of our method using various numerical simulations.

The least squares Monte Carlo method presented and discussed in this paper provides a viable alternative to price VIX derivatives in the context of a generalized framework of stochastic Volterra models. Compared to the nested Monte Carlo method, it requires much lesser computation cost to achieve the same error rate.

The choice of regression model plays a very important role in the LSMC. Ideally, it should be general enough to not require knowledge of the structure of h_T . Moreover, it should be robust to outliers, since estimates of h_T are potentially noisy. Finally, it needs to be itself not too computationally expensive, both in terms of memory and computing time. The LSMC Hermite has these characteristics but does not generalize well to higher dimensions. We explored how Random Forests and Neural Networks can handle the one-dimensional case, but more progress can be done with respect to optimizing these kinds of models. Thus, a robust and flexible method more tailored to stochastic Volterra models is a potential area of research. In particular, it would be interesting to develop a full non-Markovian method capable of efficiently handling the rough vol-of-vol case.

Finally, being able to calibrate these models is essential for their usage in industry. Directly applying a conventional calibration algorithm to the LSMC method here presented as it is, may be impractical. However, using the Deep Calibration approach of [Bayer et al., 2019], a dataset of VIX option valuations can be generated via the LSMC, and then a Neural Network can be trained to learn the mapping. This Neural Network can then be fed to a classical calibration algorithm like the Levenberg-Marquardt. The design of such network and posterior model calibration is also a problem to be studied in further research.

A Appendix - Markovian Independent vol-of-vol

In this section we present an important result proved in [Horvath et al., 2020] that allows us to calculate the infinite dimensional state variable h_T , defined in Eq. (3.8). We also provide a proof of Proposition 3.1.

In the model of [Horvath et al., 2020], the vol-of-vol Γ is assumed conservative, time-homogeneous and affine. By Theorem 2.7 and Proposition 9.1 of [Horvath et al., 2020], the infinitesimal generator of Γ can be written as

$$\mathcal{L}f(x) = k(\theta - x)\frac{\partial f}{\partial x}(x) + \frac{\delta^2 x}{2}\frac{\partial^2 f}{\partial x^2}(x) + \int_0^{+\infty} [f(x+z) - f(x)][m(dz) + x\mu(dz)]$$

for constants $k, \theta, \delta \geq 0$ and m, μ positive measures on \mathbb{R}^+ such that

$$\int_0^{+\infty} (z \wedge 1)[\mu(dz) + m(dz)] < +\infty.$$

Theorem A.1 (adapted from Prop. 3 in [Horvath et al., 2020])

Let $T > 0$ and define

$$R(u) = -ku + \frac{\delta^2}{2}u^2 + \int_0^{+\infty} (e^{zu} - 1)\mu(dz)$$

$$F(u) = k\theta u + \int_0^{+\infty} (e^{zu} - 1)m(dz)$$

Suppose that there exists $A > 0$ such that

$$\int_1^{+\infty} ze^{zA}[\mu(dz) + m(dz)] < +\infty$$

$$2G(T) + T(0 \vee R(A)) \leq A.$$

Then, the Ricatti ODE

$$\frac{\partial}{\partial t}\psi(t) = 2g^2(t) + R(\psi(t))$$

$$\frac{\partial}{\partial t}\phi(t) = F(\psi(t))$$

with initial condition $\psi(0) = \phi(0) = 0$ has a solution on $[0, T]$, where $0 \leq \psi \leq A$ and $0 \leq \phi \leq TF(A)$. Moreover

$$h_t(u, \gamma) := \mathbb{E} \left[\exp \left(2 \int_t^u g(u-s)^2 \Gamma_s ds \right) \mid \Gamma_t = \gamma \right] = \exp [\gamma \psi(u-t) + \phi(u-t)].$$

Let us now present a simple proof of Proposition 3.1.

Proof of Proposition 3.1. We start by noting that

$$\begin{aligned} \xi_t(u) &= \mathbb{E}[v_u \mid \mathcal{F}_t] \\ &= A_0(u)\mathbb{E}[E_{0,t}(u)E_{t,u}(u) \mid \mathcal{F}_t] \\ &= A_0(u)E_{0,t}(u)h_t(u) \\ &= \frac{\xi_0(u)}{h_0(u)}E_{0,t}(u)h_t(u). \end{aligned}$$

Thus

$$A_t(u) = A_0(u)E_{0,t}(u)$$

and hence

$$v_u = A_t(u)E_{t,u}(u).$$

If Γ is assumed Markovian, given Γ_t , $E_{t,u}$ is independent of \mathcal{F}_t . \square

B Appendix - Random Forests

In this section, we present some important results found in [Louppe, 2014], concerning ensemble methods in general and Random Forests in particular. When trying to learn the function h_t of Eq. (3.8), we may not know its structure in advance. Thus, we wish to use a regression model that is capable of learning a large range of functions, which often comes with added complexity. The price to pay for this, unfortunately, is overfitting. The main advantage of ensemble methods is that they allow us to reduce overfitting while at the same time not constraining the class of functions the regression can learn.

The decision tree regressor essentially approximates the target function by a step function. Since step functions are dense in L^1 , decision trees are able to approximate a wide range of functions. The main issue with decision trees is their propensity to overfit. Overfitting can be loosely defined as an extreme sensitivity to training data. More formally, we can measure it in terms of the variance of model predictions with respect to the (random) data used to train the model. In order to make a rigorous treatment of these matters, let us establish some notation.

Let X and Y be random variables in some probability space, taking values in \mathcal{X} and \mathcal{Y} , respectively. We wish to give our best prediction of Y given an observed value of X . A regression model φ is a function such that, given training data $\mathcal{L} \in \mathcal{P}(\mathcal{X} \times \mathcal{Y})$, produces a function $\varphi_{\mathcal{L}} : \mathcal{X} \rightarrow \mathcal{Y}$. To this function we can associate the random variable $\varphi_{\mathcal{L}}(X)$. The mean squared error (MSE) of the regression model given the training data \mathcal{L} can be defined as

$$MSE(\varphi | \mathcal{L}) = \mathbb{E} [(\varphi_{\mathcal{L}}(X) - Y)^2],$$

where the expectation is taken over the probability space where X and Y exist but the function $\varphi_{\mathcal{L}}$ (and hence the training data) is fixed. If $Y = f(X)$ is a deterministic function of X , the expectation is taken with respect to a deterministic random variable. In this case, it is possible that $\varphi_{\mathcal{L}}$ predicts Y perfectly (if $\varphi_{\mathcal{L}} = f$). If, on the other hand, Y is not solely determined by X , there may be a positive lower bound to $MSE(\varphi | \mathcal{L}, X = x)$. The Bayes regression model is defined as the one that achieves such lower bound, and it does not depend of any training data \mathcal{L} . It is of course given by

$$\phi_B(x) = \mathbb{E}[Y | X = x].$$

Since we usually do not know the distribution of $X \times Y$, we have no access to the Bayes regressor. Note that the Bayes regressor is optimal for each possible value x of X . Indeed, it minimizes the MSE for each fixed $X = x$:

$$MSE(\varphi | \mathcal{L}, X = x) = \mathbb{E} [(\varphi_{\mathcal{L}}(x) - Y)^2 | X = x].$$

In the above expectation, only Y is random, and it follows its conditional distribution given $X = x$.

We are interested in knowing the model error across all \mathcal{L} for each $x \in X$

$$MSE(\varphi | X = x) = \mathbb{E}[MSE(\varphi | \mathcal{L}, X = x)].$$

Here the expectation follows the distribution of all possible training sets \mathcal{L} and that of Y conditioned on $X = x$.

Theorem B.1 (Bias Variance Decomposition (thm 4.1. in [Louppe, 2014]))

Let $x \in \mathcal{X}$. Let the MSE for the Bayes model be denoted by

$$noise(x) = MSE(\varphi_B | X = x).$$

Consider the average prediction of our regression model φ across all possible training sets

$$\bar{\varphi}(x) := \mathbb{E}[\varphi_{\mathcal{L}}(x)]$$

and define the bias with respect to the Bayes model

$$bias(x) = \varphi_B(x) - \bar{\varphi}(x) = \mathbb{E}[Y - \bar{\varphi}(X) | X = x].$$

Let us also denote the variance of $\bar{\varphi}(x)$ across training sets by

$$var(x) = \mathbb{E}[(\varphi_{\mathcal{L}}(x) - \bar{\varphi}(x))^2].$$

then the MSE can be decomposed as follows

$$MSE(\varphi | X = x) = noise(x) + bias^2(x) + var(x).$$

the term $noise(x)$ is independent of the choice of regression model and can be seen as the intrinsic lower bound of the MSE posed by the problem. The Bayes regressor of course has zero bias and variance and thus attains this lower bound. For our choice of regression model, we must examine the bias-variance trade-off. A simpler model might achieve lower variance but be biased. A complex model might have small or even zero bias but might have high variance and thus be very sensible to training data. This phenomenon is what is usually known in Machine Learning jargon as overfitting.

the main idea of ensemble methods such as Random Forests, is to reduce overfitting, that is $var(x)$, while not increasing the bias. They are composed of M randomized models and their prediction is given by the average across each member of the ensemble

$$\varphi(x | \mathcal{L}, \theta) = \frac{1}{M} \sum_{i=1}^M \varphi_{\mathcal{L}}^i(x),$$

for randomized models φ^i (ensemble members). In order to be rigorous, let us say each randomized regressor φ^i is determined by the value of certain random parameter θ_i . Note that now the trained regression model depends not only on the choice of the training set \mathcal{L} but also of the choice of θ_i . thus, the expected value in the mean prediction for the ensemble method is taken across both training sets and random parameters θ

$$\bar{\varphi}(x) = \mathbb{E}[\mathbb{E}[\varphi_{\mathcal{L},\theta}(x)]] .$$

For a fixed training set \mathcal{L} and a fixed x , $\varphi_i(x)$ is a random variable, where the randomness exists exclusively through θ_i .

Theorem B.2 (thm 4.3. in [Louppe, 2014])

Assume the θ_i are i.i.d.. Thus, the variance is equal across each member of the ensemble:

$$\sigma^2(x) = var_{\varphi_i}(x) \forall i = 1, \dots, M.$$

By linearity of the expectation and the tower property, the bias of the ensemble is the same as the bias of any individual member

$$bias(x) = \varphi_B(x) - \bar{\varphi}(x) = \varphi_B(x) - \bar{\varphi}^i(x).$$

Let $\rho(x)$ denote the correlation coefficient between $\varphi^i(x)$ and $\varphi^j(x)$. These two random variables only depend on θ_i and θ_j , which are assumed i.i.d.. Thus, $\rho(x)$ does not depend on i and j . Then $\rho(x) \geq 0$ and the variance of the ensemble is given by

$$var(x) = \sigma^2(x) \left(\rho(x) + \frac{1 - \rho(x)}{M} \right).$$

In particular, $var(x) \leq \sigma^2(x)$.

B APPENDIX - RANDOM FORESTS

Thus, the stronger the random effects in the ensemble (meaning $\rho(x) \rightarrow 0$), the greater is the variance reduction and hence the lower the MSE of the ensemble method.

We conclude by noting that Random Forests do not totally solve the problem of overfitting but are an improvement when compared to decision trees. The improvement is more evident in higher dimensional problems with many predictors, since there it is possible to randomize the predictors across each different tree in the forest. This should inform our analysis of the Random Forest performance in lower dimensional problems such as the one presented in this paper.

References

- [Alòs et al., 2018] Alòs, E., García-Lorite, D., and Muguruza, A. (2018). On smile properties of volatility derivatives and exotic products: understanding the VIX skew. Preprint arXiv:1808.03610.
- [Alòs et al., 2007] Alòs, E., Leon, J., and Vives, J. (2007). On the short-time behavior of the implied volatility for jump-diffusion models with stochastic volatility. *Finance and Stochastics*, 11:571–589.
- [Barndorff-Nielsen and Schmiegel, 2009] Barndorff-Nielsen, O. E. and Schmiegel, J. (2009). Brownian semistationary processes and volatility/intermittency. In Albrecher, H., Runggaldier, W. J., and Schachermayer, W., editors, *Advanced Financial Modelling*, pages 1–26. De Gruyter.
- [Bayer et al., 2016] Bayer, C., Friz, P., and Gatheral, J. (2016). Pricing under rough volatility. *Quantitative Finance*, 16(6):887–904.
- [Bayer et al., 2019] Bayer, C., Horvath, B., Muguruza, A., Stemper, B., and Tomas, M. (2019). On deep calibration of (rough) stochastic volatility models. Preprint arXiv:1908.08806.
- [Bennedsen et al., 2017] Bennedsen, M., Lunde, A., and Pakkanen, M. S. (2017). Hybrid scheme for Brownian semistationary processes. *Finance and Stochastics*, 21(4):931–965.
- [Chicago Board Options Exchange, 2019] Chicago Board Options Exchange (2019). VIX: CBOE volatility index.
- [Da Fonseca and Zhang, 2019] Da Fonseca, J. and Zhang, W. (2019). Volatility of volatility is (also) rough. *Journal of Futures Markets*, 39(5):600–611.
- [El Euch et al., 2018] El Euch, O., Fukasawa, M., and Rosenbaum, M. (2018). The microstructural foundations of leverage effect and rough volatility. *Finance and Stochastics*, 22(2):241–280.
- [Euch and Rosenbaum, 2017] Euch, O. and Rosenbaum, M. (2017). Perfect hedging in rough heston models. *Annals of Applied Probability*, 28:241–280.
- [Fukasawa, 2015] Fukasawa, M. (2015). Short-time at-the-money skew and rough fractional volatility. *Quantitative Finance*, 17:189–198.
- [Gatheral et al., 2018] Gatheral, J., Jaisson, T., and Rosenbaum, M. (2018). Volatility is rough. *Quantitative Finance*, 18(6):933–949.
- [Gatheral et al., 2020] Gatheral, J., Jusselin, P., and Rosenbaum, M. (2020). The quadratic rough Heston model and the joint S&P 500/VIX smile calibration problem. Preprint available at SSRN: <https://ssrn.com/abstract=3514894> or <http://dx.doi.org/10.2139/ssrn.3514894>.
- [Goutte et al., 2017] Goutte, S., Ismail, A., and Pham, H. (2017). Regime-switching stochastic volatility model: estimation and calibration to vix options. *Applied Mathematical Finance*, 24(1):38–75.
- [Guo and Loeper, 2018] Guo, I. and Loeper, G. (2018). Pricing bounds for volatility derivatives via duality and least squares monte carlo. *Journal of Optimization Theory and Applications*, 179(2):598–617.
- [Guyon, 2019] Guyon, J. (2019). The joint S&P 500/VIX smile calibration puzzle solved. Risk, April 2020, Available at SSRN: <https://ssrn.com/abstract=3397382> or <http://dx.doi.org/10.2139/ssrn.3397382>.
- [Horvath et al., 2020] Horvath, B., Jacquier, A., and Tankov, P. (2020). Volatility options in rough volatility models. *SIAM Journal on Financial Mathematics*, 11:437–469.

REFERENCES

- [Horvath et al., 2019] Horvath, B., Muguruza, A., and Tomas, M. (2019). Deep learning volatility. *SSRN Electronic Journal*.
- [Lelong, 2020] Lelong, J. (2020). Pricing path-dependent Bermudan options using wiener chaos expansion: an embarrassingly parallel approach. *Journal of Computational Finance*, 24(2):1–31.
- [Longstaff and Schwartz, 2001] Longstaff, F. and Schwartz, E. (2001). Valuing American options by simulation: A simple least-squares approach. *Review of Financial Studies*, 14:113–47.
- [Louppe, 2014] Louppe, G. (2014). *Understanding Random Forests: From Theory to Practice*. PhD thesis, University of Liège.
- [McCrickerd and Pakkanen, 2018] McCrickerd, R. and Pakkanen, M. S. (2018). Turbocharging monte carlo pricing for the rough Bergomi model. *Quantitative Finance*, 18(11):1877–1886.
- [Stone, 2019] Stone, H. (2019). Calibrating rough volatility models: a convolutional neural network approach. *Quantitative Finance*, 20:1–14.