



**HAL**  
open science

## Path Shadowing Monte-Carlo

Rudy Morel, Stéphane Mallat, Jean-Philippe Bouchaud

► **To cite this version:**

Rudy Morel, Stéphane Mallat, Jean-Philippe Bouchaud. Path Shadowing Monte-Carlo. 2023. hal-04177835

**HAL Id: hal-04177835**

**<https://hal.science/hal-04177835>**

Preprint submitted on 6 Aug 2023

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Path Shadowing Monte-Carlo

Rudy Morel<sup>a</sup>, Stéphane Mallat<sup>a,b</sup> & Jean-Philippe Bouchaud<sup>c</sup>

<sup>a</sup>*École Normale Supérieure, Paris, France*

<sup>b</sup>*Collège de France, Paris, France*

<sup>c</sup>*Capital Fund Management, Paris, France*

**Abstract**—We introduce a *Path Shadowing Monte-Carlo* method, which provides prediction of future paths, given any generative model. At any given date, it averages future quantities over generated price paths whose past history matches, or ‘shadows’, the actual (observed) history. We test our approach using paths generated from a maximum entropy model of financial prices, based on a recently proposed multi-scale analogue of the standard skewness and kurtosis called ‘Scattering Spectra’ [1]. This model promotes diversity of generated paths while reproducing the main statistical properties of financial prices, including stylized facts on volatility roughness. Our method yields state-of-the-art predictions for future realized volatility and allows one to determine conditional option smiles for the S&P500 that outperform both the current version of the Path-Dependent Volatility model and the option market itself<sup>1</sup>.

**Index Terms**—volatility prediction, option pricing, wavelets

## I. INTRODUCTION

Modelling future price scenarios is crucial for risk control, for pricing and hedging contingent claims (like options), and, possibly, for detecting arbitrage opportunities. Recently, machine learning auto-regressive models [2]–[4] manage to learn from data the distribution  $p(x|x_{\text{past}})$  of log-prices  $x$  conditioned on past history  $x_{\text{past}}$ . When trained with a prediction loss, such models generally achieve excellent prediction results. However, their training requires very large amount of data which are usually not available for financial prices.

On the other hand, low-parameterized generative models, i.e. models  $p_{\theta}(x)$  of  $p(x)$  with few parameters  $\theta$ , have been extensively studied in the financial literature [5]–[9]. However, two main challenges come to the fore. First, these models may not reproduce some important statistics of real financial prices due to oversimplified or flawed assumptions, or due to the fact that they are calibrated on external data such as observed option smiles. Second, it may not be straightforward to condition these models on the realized past at a specific date, in other words, obtaining a model of  $p(x|x_{\text{past}})$ . Whereas conditioning is eased by considering Markovian models with a small number of factors [5], [9], such a strong assumption is often much too simplistic.

In this paper, we attempt to address both challenges. Our main contribution is to introduce a new method, that we call *Path Shadowing Monte-Carlo* (PS-MC), which can be used within any generative model of  $p(x)$  to yield a model of  $p(x|x_{\text{past}})$ . Our approach for modelling the distribution  $p(x)$ ,




Fig. 1: Path shadowing. Given current past history  $\tilde{x}_{\text{past}}$  (red), we scan for paths  $x$  (gray) in a generated dataset whose past history satisfies  $x_{\text{past}} \approx \tilde{x}_{\text{past}}$ . Such paths  $x$  are said to *shadow*  $\tilde{x}_{\text{past}}$ , they provide insights on the future. Predictions are obtained through Monte-Carlo on such shadowing paths.

summarized in section II, is to define a minimal set of statistics describing financial prices that should be reproduced by the generating process. Such statistics should be well-estimated on limited data, while specifying ‘relevant’ properties of the process, in a sense made precise below. This bias-variance tradeoff was addressed in our previous work [1] in the general case of multi-scale processes, where it was shown that a good description of financial prices can be achieved by multi-scale analogues of the classical skewness and kurtosis, called *Scattering Spectra*, which is motivated in section II.

A model based on these statistics captures all important stylized facts such as fat-tail distributions, intermittency, leverage effect and the ‘Zumbach effect’ [1]. Section III characterizes the average shape of option smiles generated by our Scattering Spectra model and shows that it accounts in particular for the power-law behaviour of the at-the-money skew as a function of maturity, which were recently argued to be a specific feature of rough volatility models [7], [10]. The power-law behaviour of the kurtosis, first reported in [11], is also remarkably well accounted for.

‘Path shadowing’ is presented in section IV. It consists in softening the conditioning on a given past history  $x_{\text{past}}$ . In a nutshell, it amounts to scanning a large generated dataset, in search of paths whose history closely ‘shadows’ the actual history (see Fig. 1 for an illustration). A Path Shadowing Monte-Carlo method then averages the quantity of interest over the future of such matching paths. The term ‘shadowing’ is freely inspired by the shadowing principle in chaotic dynamical systems [12]–[15]. Intuitively, it states that a path which is uniformly close to a true orbit will stay close (shadow) a

<sup>1</sup>This work is supported by the PRAIRIE 3IA Institute of the French ANR-19-P3IA-0001 program and the ENS-CFM models and data science chair.

true path for all time.

This method can effectively be seen as a kernel method, with a causal path embedding to reduce the dimensionality of recent past history. Unlike other recent kernel methods, such as signature kernels [16], [17] that rely on a low-parametric model for  $p(x|x_{\text{past}})$ , Path Shadowing Monte-Carlo relies solely on a model of  $p(x)$ . It thus circumvents the exact conditioning of a generative model to a given past history  $x_{\text{past}}$ . Its performance depends directly on the accuracy of this generative model and its ability to produce a variety of paths with correct statistical dependencies. Section IV-C shows that when performed with our maximum entropy Scattering Spectra model of financial prices, PS-MC yields state-of-the-art volatility prediction.

Section V uses Path Shadowing Monte-Carlo for obtaining *conditional* option smiles (i.e. option prices at a given date) through Hedged Monte-Carlo with shadowing paths. By construction, such smiles depend only on the log-price process distribution  $p(x)$  and provide a counterpart to smiles obtained from option market data. A ‘trading game’ then allows us to show that our option smiles correctly anticipate non-trivial future price movements, and compares favourably with state of the art models such as the Path-Dependent Volatility model (PDV) of ref. [9]. Codes for both our generative model and Path Shadowing Monte-Carlo are available at <https://github.com/RudyMorel/shadowing>.

## II. A MULTI-SCALE STATISTICAL MODEL FOR FINANCIAL PRICES

Statistical models of financial prices aim at reproducing statistics of the price process only. Price time series exhibit numerous non-Gaussian features, which are difficult to capture within standard low-parametric models, whose number of parameters have been incrementally increased in the literature over the past decades, see e.g. [5]–[7], [9], [18]. An alternative route is to define a set of characteristic statistics of (log-)prices and impose that they should be accurately reproduced by the model. We denote as  $\Phi(x)$  such statistics, for example the empirical mean and variance of log-returns  $\Phi(x) = (\langle \delta x(t) \rangle_t, \langle |\delta x(t)|^2 \rangle_t)$  where, throughout this paper,  $\langle \dots \rangle_t$  denotes empirical averages over time  $t$ . Section II-A presents maximum entropy models that allow defining models from a given vector of statistics  $\Phi(x)$ . In the simple case of mean and variance, the maximum entropy model coincides with the Gaussian random walk.

The set  $\Phi(x)$  must be chosen carefully and is governed by a bias-variance tradeoff. It should contain enough relevant statistics of prices for the model to be realistic and accurate – reducing the model bias. However, such statistics must be well estimated on the single available historical realization of  $x$  – reducing the model variance. The construction of a set  $\Phi$  that meets these requirements, called the *Scattering Spectra* (SS), was proposed in [1] in the general case of multi-scale processes, which fortunately includes financial data [6], [19], [20].

We show in section II-B that such statistics correspond to low moment multi-scale analogue of the classical skewness

and kurtosis of log-returns. We show that even the most recent low-dimensional parametric models fail to accurately account for these statistics. Such discrepancies turn out to be highly relevant when one wants to predict future realized volatility and option smiles, and highlights the limitations of traditional models, which our approach allows one to overcome.

In [1], we have shown that a Scattering Spectra model properly captures the main properties of financial log-returns, in particular of the S&P500 (the US major stock index). In the following, we show that it also quantitatively reproduces the average behavior of option smiles of different maturities, in particular the maturity-dependent skewness that reflects volatility roughness [10] and the so-called skew-stickiness ratio [21], [22].

### A. Maximum Entropy Models

We denote as  $\tilde{x} \in \mathbb{R}^N$  the observed historical realization of log-prices over  $N$  days. Given a vector of  $d$  statistics  $\Phi(\tilde{x}) \in \mathbb{R}^d$  estimated on  $\tilde{x}$ , a maximum entropy model  $p_\theta$  with moment constraint  $\mathbb{E}_{p_\theta}\{\Phi(x)\} = \Phi(\tilde{x})$ , if it exists, has an exponential probability distribution [23]

$$p_\theta(x) = Z_\theta^{-1} e^{-\langle \theta, \Phi(x) \rangle}. \quad (1)$$

for certain  $\theta \in \mathbb{R}^d$ . Estimating the parameters  $\theta$  of model (1) is computationally expensive, in particular when the number of statistics  $d$  is large [24], [25]. To avoid this issue, we consider in this paper microcanonical maximum entropy models which approximate model (1). These models, together with a sampling algorithm are described in Appendix A.

Maximum entropy models depend only on the vector of statistics  $\Phi(x)$ . The model accuracy can be improved by enriching the set  $\Phi(x)$ . However, we must take into account the problem of estimating  $\Phi(x)$  from the single realization of the process  $\tilde{x}$ . The SS model imposes  $\mathbb{E}_{p_\theta}\{\Phi(x)\} = \Phi(\tilde{x})$ , thus for  $p_\theta$  to be a good approximation of the true distribution  $p$ , one needs  $\Phi(\tilde{x})$  to be close to the true  $\mathbb{E}_p\{\Phi(x)\}$ . This amounts to having low-variance statistics  $\Phi$ . A good choice of  $\Phi$  is presented in the next section.

### B. The Scattering Spectra (SS)

A standard way of characterizing the price process is through their trend, volatility, skewness and kurtosis. These are obtained from moments of order 1, 2, 3 and 4 on log-returns

$$\mathbb{E}\{\delta x(t)\}, \mathbb{E}\{\delta x(t)^2\}, \mathbb{E}\{\delta x(t)^3\}, \mathbb{E}\{\delta x(t)^4\}$$

However such moments do not characterize the time-structure of log-returns, but rather their one-point distribution. One could consider the same moments on multi-scale increments

$$\delta_\ell x(t) = x(t) - x(t - \ell) \quad (2)$$

for different lags  $\ell$ , but we still obtain a poor description of  $x$ . For example, these moments do not pick up time-asymmetry, since changing  $\delta x(t)$  into  $\delta x(-t)$  leaves these moments unchanged. Another disadvantage of multi-scale increments (2) is that they exhibit as many scales  $1 \leq \ell \leq N$  as the number

of days  $N$ , which seems redundant, specially in view of the known scale-invariant properties of  $x$ .

The *Scattering Spectra* introduced in [1] capture the main non-Gaussian properties of financial prices: fat tailed log-return distributions, sign-asymmetry, time-asymmetry, volatility clustering and volatility roughness. It consists of  $d = \mathcal{O}(\log_2^3 N)$  statistics only, that are low-order moments (order 1 and 2 only) and can thus be accurately estimated on the historical realization  $\tilde{x}$  of size  $N$ .

We present here the main steps for building such  $\Phi$  and we refer the reader to [1] for more details about the construction.

### Step 1. Wavelet increments.

Log-prices variation have interesting structure at all scales. However, it is not necessary to consider all scales  $\ell$  in (2) to characterize them efficiently. Standard increments  $\delta_\ell x(t)$  (2) are obtained by convolution of  $x$  with the filter  $g_\ell = \delta_0 - \delta_\ell$ . Wavelet increments replace  $g_\ell$  by wavelet filters  $\psi_j$  obtained by dilation of a regular mother wavelet  $\psi$

$$W_j x(t) = x \star \psi_j(t) \quad \text{where} \quad \psi_j(t) = 2^{-j} \psi(2^{-j} t). \quad (3)$$

The mother wavelet  $\psi$  has a zero average  $\int \psi(t) dt = 0$  and its Fourier transform  $\widehat{\psi}(\omega) = \int \psi(t) e^{-i\omega t} dt$ , which is real, is mostly concentrated at frequencies  $\omega \in [\pi, 2\pi]$ . All numerical calculations in this paper are performed with a complex Battle-Lemarié wavelet [26], [27]. Fig. 7 shows the real and imaginary parts of  $\psi$  as well as its Fourier transform. We refer the reader to Appendix B for more properties.

Analogous to (2), wavelet increments (3) can be seen as multi-scale increments at scales  $\ell = 2^j$  with  $j = 1, \dots, J$ . However, scales are now defined as bins of frequencies  $[2^{-j}\pi, 2^{-j+1}\pi]$  corresponding to the supports of wavelet filters  $\psi_j$ . The largest scale  $2^J$  is chosen to be smaller than the size  $N$  of  $\tilde{x}$ . This yields at most  $\log_2 N$  scales instead of  $N$  lags  $\ell$ .

Histograms of generalized increments  $W_j x$  can be constrained by order 1 and order 2 moments  $\mathbb{E}\{|W_j x(t)|\}$ ,  $\mathbb{E}\{|W_j x(t)|^2\}$  which are estimated through empirical averages. The quantity

$$\Phi_1(x)[j] = \frac{\langle |W_j x(t)| \rangle_t}{\langle |W_j x(t)|^2 \rangle_t} \quad (4)$$

is a low-moment measure of kurtosis. Compared to its order 4 counterpart, it is less sensitive to large values. The more peaked at zero the distribution, the smaller the value of  $\Phi_1(x)$  and the higher the kurtosis [28]. The order 2 moment is

$$\Phi_2(x)[j] = \langle |W_j x(t)|^2 \rangle_t \quad (5)$$

and quantifies the average volatility at scale  $2^j$  on the period.

### Step 2. Time-scale dependencies.

Multi-scale increments  $W_j x(t)$  are indexed by time  $t$  and scale  $2^j$ . Such map exhibits dependencies across time and scales that are crucial to characterize the distribution of financial prices. For example, volatility clustering is attested by the fact that  $W_j x(t)$  has long-range time correlations. Beyond this well-known stylized fact, the authors of [1] have shown that scale dependencies are crucial to fully characterize the




Fig. 2: Standard statistics of log-returns in the Scattering Spectra (SS) model (orange) compared to S&P observed data (blue). Top graphs: time series of the S&P and generated by the model. Bottom graphs: (a) Histogram of daily log-returns  $\delta x$ . (b) Structure functions  $\langle |\delta_\ell x(t)|^q \rangle_t$ . (c) Leverage correlation  $\langle \delta x(t - \tau) |\delta x(t)|^2 \rangle_t$  on normalized increments.

non-Gaussian nature of time series. Natural descriptors for such scale dependencies are order 2, 3 and 4 moments

$$\mathbb{E}\{Wx Wx^*\}, \quad \mathbb{E}\{Wx |Wx|^2\}, \quad \mathbb{E}\{|Wx|^2 |Wx|^2\}$$

where the products are taken across times  $t, t'$  and scales  $j, j'$ . In practice, estimating order 3 and order 4 moments is very difficult because of the variance induced by large events. In order to circumvent this problem, we replace  $|Wx|^2$  by  $|Wx|$  and define the following non-linear correlations of wavelet increments

$$\mathbb{E}\{Wx Wx^*\}, \quad \mathbb{E}\{Wx |Wx|\}, \quad \mathbb{E}\{|Wx| |Wx|\} \quad (6)$$

Owing to the compression properties of wavelets, the first matrix  $\mathbb{E}\{Wx Wx^*\}$  is quasi-diagonal and its diagonal coefficients are already estimated by (5), see [1].

### Step 3. Low-moment multi-scale skewness and kurtosis.

Just like for standard skewness and kurtosis that are normalized moments, we normalize the second and third matrices  $\mathbb{E}\{Wx |Wx|\}$  and  $\mathbb{E}\{|Wx| |Wx|\}$  in (6) by  $\mathbb{E}\{|Wx|^2\}$ . One can show that the only non-negligible coefficients in the third matrix are obtained for  $t = t'$  and  $j \geq j'$ , and are estimated through

$$\Phi_3(x)[j, j'] = \frac{\langle W_j x(t) |W_{j'} x(t)| \rangle_t}{\langle |W_j x(t)|^2 \rangle_t^{\frac{1}{2}} \langle |W_{j'} x(t)|^2 \rangle_t^{\frac{1}{2}}}. \quad (7)$$

These are analogous to the standard low-moment skewness  $\mathbb{E}\{Y|Y|\}$  of a normalized random variable  $Y$  with  $\mathbb{E}\{Y^2\} = 1$ . Other than sign-asymmetry, these complex coefficients also measure time-asymmetry through their phase. Indeed, if log-returns are time-reversible  $\delta x(-t) \stackrel{d}{=} \delta x(t)$  then  $\text{Im} \Phi_3(x) = 0$ . One typical example is the leverage asymmetric correlation.

The fourth matrix  $\mathbb{E}\{|Wx| |Wx|\}$  in (6) contains kurtosis information. If  $x$  is Gaussian, then for different scales

$j \neq j'$  the Gaussian processes  $W_j x$  and  $W_{j'} x$  are decorrelated, thus independent. It follows that  $\mathbb{E}\{|W_j x| | W_{j'} x|\} = \mathbb{E}\{|W_j x|\} \mathbb{E}\{|W_{j'} x|\}$  and these coefficients boil down to the low-moment kurtosis (4). For the log-price process  $x$ , these coefficients capture long-range non-Gaussian correlation between volatility at different scales  $j, j'$  and different times  $t, t'$ .

However, the matrix  $\mathbb{E}\{|W x| | W x|\}$  contains too many coefficients to be accurately estimated on a single realization  $\tilde{x}$ . We again rely on compression properties of wavelets to approximate such matrix by cascading a second wavelet operator  $W$ , which yields a quasi-diagonal matrix  $\mathbb{E}\{W | W x| W | W x|^*\}$  where we define generalized increments of volatility as

$$W_{j_2} |W_{j_1} x|(t) = |x \star \psi_{j_1}| \star \psi_{j_2}(t).$$

The non-negligible diagonal coefficients are estimated through an empirical average which yields for  $j_1 \leq j'_1 < j_2$

$$\Phi_4(x)[j_1, j'_1, j_2] = \frac{\langle W_{j_2} |W_{j_1} x|(t) W_{j_2} |W_{j'_1} x|(t) \rangle_t}{\langle |W_{j_1} x(t)|^2 \rangle_t^{\frac{1}{2}} \langle |W_{j'_1} x(t)|^2 \rangle_t^{\frac{1}{2}}}. \quad (8)$$

These are analogous to the classical low-moment kurtosis. These complex coefficients also capture time-asymmetry through their complex phase. If the log-return process  $\delta x$  is time-reversible then  $\text{Im} \Phi_4(x) = 0$ .

We therefore define our Scattering Spectra  $\Phi$  as the collection of (i) estimated average volatility (5), (ii) multi-scale skewness (7) and (iii) multi-scale kurtosis (4,8)

$$\Phi(x) = (\Phi_1(x), \Phi_2(x), \Phi_3(x), \Phi_4(x)). \quad (9)$$

In total,  $\Phi$  consists of  $\mathcal{O}(\log_2^3(N))$  order 1 and order 2 statistics for a trajectory of size  $N$  and can be estimated with low-variance.

Note that  $\Phi$  does not rely explicitly on the one-point distribution of increments  $\delta_\ell x(t)$ . Numerical experiments have shown that slight discrepancies may appear, in particular in order 0 moments  $\mathbb{P}(\delta_\ell x(t) > 0)$  which explicitly appear in low-moment smile expansions [28]. We thus complement  $\Phi_3(x)$  with the moments

$$\mathbb{P}(\delta_\ell x(t) > 0)$$

for  $\ell = 2^j, j = 1, \dots, J$ , that are estimated through empirical averages  $\langle \text{sigmoid}(\delta_\ell x(t)) \rangle_t$  where  $\text{sigmoid}(x) = (1 + e^{-x})^{-1}$ . This adds very few coefficients to our scattering spectra  $\Phi(x)$ .

The Scattering Spectra (9) thus provide an enriched set of statistics that can be used to quantify model error and interpret any discrepancy. As an example, we revisit through this lens the, low-parametric, Path-Dependent Volatility (PDV) model introduced by Guyon & Lekeufack [9]. Although more parsimonious in terms of number of parameters, several stylized facts are in fact not accurately reproduced by such a model, see a more precise discussion in Appendix D, Fig. 11.

Based on the Scattering Spectra  $\Phi$ , we have at our disposal a statistical model of financial prices that can be used to generate faithful synthetic time series (see section II-A). For the S&P time series  $\tilde{x}$  of size  $N = 5827$  days, the Scattering Spectra model (SS model) contains  $248 \approx N/20$  real coefficients, which is the dimension of  $\Phi(x)$ . Log-return trajectories  $\delta x$

generated from the SS model are shown in Fig. 2. Validation of the SS model can be achieved by measuring observables not included in our set  $\Phi(x)$  and checking whether or not they are correctly reproduced. Standard statistics such as fat tails, volatility clustering, leverage effect and structure functions, were indeed shown to be captured by the model [1]. These are reproduced in Fig. 2. While  $\Phi(x)$  is composed of order 1 and order 2 moments only, the SS model accurately accounts for up to order 5 moments, which is quite remarkable. Another way to describe the multi-scale statistical properties of price time series is through maturity dependent option smiles, which we discuss in the next section.

### III. THE AVERAGE SMILE AS AN ALTERNATIVE STATISTICAL CHARACTERIZATION

In this section we validate the SS model by considering historical option pricing as an alternative, intuitive way to characterize the multi-scale, non-Gaussian statistics of price time series. The *average smile* is the unconditional option smile obtained by pricing hedged options using all historical snippets of prices of length equal to the maturity of the option [28], [29]. Even if real option smiles must be conditioned on a specific past price path [9] and are therefore almost never equal to the *average smile*, its shape reveals some interesting, non-trivial properties of prices time series, such as volatility ‘roughness’ (see below).

Option pricing is performed through the Hedged Monte-Carlo method [29], that converts historical probabilities (either real or synthetic) into ‘risk-neutral’ ones. Options are hedged daily, with zero interest rate, on the 6000 price snippets of lengths 150 days available from 2000 to 2023 all rescaled such that the initial price is 100. The average implied volatilities  $\sigma(T, K)$  are obtained from option prices  $\mathcal{C}(T, K)$ . Fig. 3 compares, for different maturities  $T$ , the *average smiles* using observed S&P data and those generated with the SS model. We see that the model indeed reproduces the overall shape of the smile very well. Intuitively, the level of the average smile, its asymmetry, its concavity and its term structure are captured by  $\Phi_2$  (5),  $\Phi_3$  (7) and  $\Phi_1$  and  $\Phi_4$  (4,8)<sup>2</sup>. We have also compared the S&P average smiles with the recent Path-Dependent Volatility model of [9], calibrated such as to reproduce the same SS as best as possible – see Appendix D for more details. As a general comment, the PDV model underestimates the kurtosis of the process and correspondingly fail to capture accurately the right wing of the average smile, see Fig. 3c.

We now turn to a more refined analysis of the slope and curvature of these average smiles. We denote as  $\sigma_{\text{ATM}}(T) = \sigma(T, 100)$  the at-the-money volatility and

$$\mathcal{M} := \frac{\ln(\frac{K}{100})}{\sigma_{\text{ATM}} \sqrt{T}}$$

<sup>2</sup>Appendix C shows in more details the parameterization of the model by studying the sensitivity of the smile to the Scattering Spectra statistics  $\Phi(x)$ .

the *rescaled* log-moneyness. The slope  $\mathcal{S}_T$  and curvature  $\kappa_T$  of a smile at maturity  $T$  are defined by the order 2 expansion around the moneyness  $\mathcal{M} = 0$

$$\sigma(\mathcal{M}, T) := \sigma_{\text{ATM}}(T) \left( 1 + \mathcal{S}_T \mathcal{M} + \kappa_T \mathcal{M}^2 + o(\mathcal{M}^2) \right)$$

In the literature, it is customary to define the ATM skew  $\text{Skew}_T$  as the slope of the smile as a function of *unscaled* log-moneyness, i.e.  $\text{Skew}_T := \mathcal{S}_T / \sqrt{T}$ . For most stochastic volatility models, such skew is found to be regular when  $T \rightarrow 0$ , whereas rough volatility models predict a singular behavior  $\text{Skew}_T \propto T^{H-1/2}$  where  $H$  is the Hurst exponent of volatility, argued to be small,  $H \approx 0.1$  [7], [30], [31].

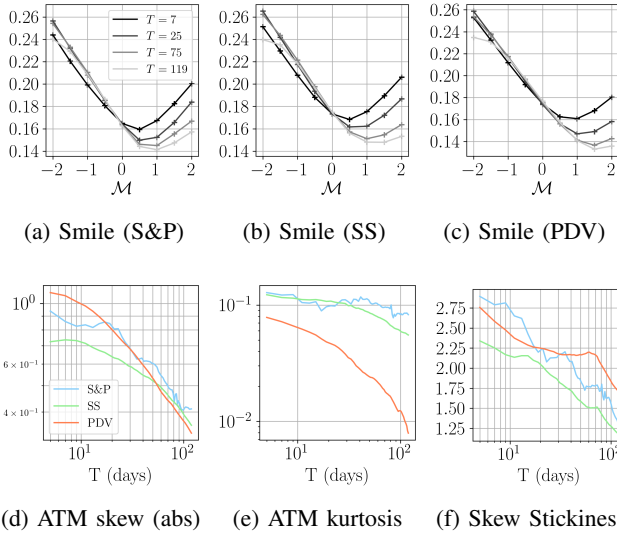


Fig. 3: Average option smiles estimated using S&P price data (a), in our Scattering Spectra (SS) model (b) and in the Path-Dependent Volatility (PDV) model (c). The SS model qualitatively captures the two regimes of the ATM skew as a function of maturity (d), with a cross-over around 20 days [18], [32], as well as the very slow decaying power-law of the ATM kurtosis [11] (e) and the behavior of the skew-stickiness ratio (f). The PDV model, on the other hand, correctly captures skewness effects (d) but fails to capture the amplitude and term structure of the ATM kurtosis (e).

Fig. 3d shows the absolute value of the ATM skew of the average smile for different maturities. In [18], [32] it was shown using option market data that the ATM skew exhibits two power-law regimes pertaining to short and long maturities, with a cutoff between the two regimes around 20 days. Interestingly, we also observe this behavior on the average smile of the S&P, which only depends on the price process, with no reference to actual option markets. Both the PDV model and the SS model track rather well such a behavior. The scaling of log-volatility increments characteristic of rough volatility models [7] or multifractal models [6] is in fact encoded in the model through the coefficients  $\mathbb{E}\{|W_{j_2}|W_{j_1}x|(t)|^2\}$  included in  $\Phi_4(x)$  (8). These consider instantaneous volatility  $|W_{j_1}x|$  at scale  $2^{j_1}$  and its increments at scales  $2^{j_2}$ .

The SS model furthermore captures two more subtle stylized facts of option smiles:

- The ATM curvature  $\kappa_T$ , related to a low moment kurtosis [28], is well captured and behaves as a slow power-law of  $T$  (first reported in [11]), both for the S&P and within the SS model. The PDV model, on the other hand, strongly underestimates  $\kappa_T$  (see Fig. 3e).
- The instantaneous change of ATM volatility  $\sigma_{\text{ATM}}(T)$  induced by a change in underlying price can be linearly regressed on  $\delta x$ . This defines the skew-stickiness ratio  $R_T$  through the following regression [21], [22]

$$\delta\sigma_{\text{ATM}}(T) = -R_T \text{Skew}_T \times \delta x + \epsilon$$

As shown in [22],  $R_T$  has a non-trivial dependence on maturity. Fig. 3f shows that both the PDV model and the SS model again reproduce quite well such a dependence.

#### IV. PATH SHADOWING MONTE-CARLO & VOLATILITY PREDICTIONS

The *average smile* exercise of the previous section is interesting insofar as it allows one to test the ability of various models to capture the distribution of price changes over different maturities. However, it fails to inform us on the power of the model to actually predict, at a given date, the distribution of price changes forward in time. Of course, this is what Finance is all about and we now introduce a framework to do precisely that.

We first assume that the real world is at least approximately stationary, in the sense that it can be approximated by a statistical model with fixed, time-independent parameters. Of course, this can only be true if the model is rich enough to generate intermittent time series that superficially appear non-stationary – such as the ones shown in Fig. 2, with alternating periods of high and low volatility that are actually described by the *same* model.

If this is the case, then given the past history  $\tilde{x}_{\text{past}}$  at current time  $t$ , a model-free method for predicting the unknown future  $\tilde{x}_{\text{future}}$  is to look for occurrences similar to  $\tilde{x}_{\text{past}}$  in the historical realization of log-prices. If such occurrences can be found, what happened thereafter provides some information about the unknown future  $\tilde{x}_{\text{future}}$  at the current time  $t$ .

Finding exact occurrences of course happens with vanishing probability. We therefore introduce an embedding  $h(\tilde{x}_{\text{past}})$  that reduces the dimensionality of past trajectories and define *shadowing paths* as paths  $x$  whose past history  $h(x_{\text{past}})$  is close to  $h(\tilde{x}_{\text{past}})$  in a certain sense. Furthermore, instead of scanning the historical past, we propose in this section to scan a long dataset of paths generated using the model presented in section II.

These shadowing paths are then used as inputs of our proposed *Path Shadowing Monte-Carlo* (PS-MC) method, which allows us to obtain state-of-the-art predictions for future realized volatility. The method will be extended in the next section V to option pricing.

##### A. The Path Shadowing Monte-Carlo Method

First, at a given time  $t$ , we separate a log-price path  $x \in \mathbb{R}^N$  into its past  $x_{\text{past}} = (x(u), u \leq t)$  and future  $x_{\text{future}} = (x(u), u > t)$

$$x = (x_{\text{past}}, x_{\text{future}})$$

Let  $q(x_{\text{future}})$  be a quantity we want to predict, for example the average realized variance in the next  $T$  days  $q(x_{\text{future}}) := \sum_{u=t+1}^{t+T} |\delta x_u|^2 / T$ . In the following section, we write  $q(x) = q(x_{\text{future}})$  for simplicity. An optimal prediction of  $q(x)$  for the mean square error as a function of the observed past  $\tilde{x}_{\text{past}}$  is given by the conditional expectation

$$\mathbb{E}\{q(x) \mid x_{\text{past}} = \tilde{x}_{\text{past}}\} \quad (10)$$

with  $\mathbb{E}$  the expectation under the true distribution  $p(x)$  of log-prices. The goal is to estimate such conditional expectation.

Let us for a moment omit the conditioning on the past. The standard Monte-Carlo method estimates expectations using a finite number of realizations  $x^1, \dots, x^n$  drawn from  $p(x)$  as

$$\frac{1}{n} \sum_{k=1}^n q(x^k)$$

which converges to  $\mathbb{E}\{q(x)\}$  as  $n \rightarrow +\infty$  under independence of  $x_1, \dots, x_n$  and integrability of  $q(x)$ .

In theory, such method could apply to estimate (10), however it would require finding paths  $x^k$  such that  $x_{\text{past}}^k = \tilde{x}_{\text{past}}$ , which is all but impossible for data of finite size.

To tackle this problem, we relax strict conditioning on  $\tilde{x}_{\text{past}}$  and consider paths  $x$  whose past  $x_{\text{past}}$  is *close* to  $\tilde{x}_{\text{past}}$  in a certain sense. In order to account for possible long-range dependencies, we would like to consider a long past  $\tilde{x}_{\text{past}}$ . However, finding paths  $x_{\text{past}}$  at a given distance from  $\tilde{x}_{\text{past}}$  becomes exponentially difficult as the size of the path grows – this is the curse of dimensionality. In order to control the dimension, we consider a path embedding  $h(x_{\text{past}}) \in \mathbb{R}^M$ . Given a threshold  $\eta > 0$  we define the set of  $\eta$ -shadowing paths as

$$H_\eta(\tilde{x}_{\text{past}}) = \{x \in \mathbb{R}^N \mid \|h(x_{\text{past}}) - h(\tilde{x}_{\text{past}})\| < \eta\} \quad (11)$$

For example,  $h(x) = \delta x$  considers past log-returns, and hence log-price paths up to an additive constant. Our choice of  $h$  is detailed in the next section. The term *shadowing* is freely inspired by the shadowing principle in chaotic dynamical systems [12]–[15]. The idea is that for a certain small  $\eta$ , paths in  $H_\eta(\tilde{x}_{\text{past}})$  can be considered as true realizations of the process that can be used to faithfully compute observables.

*Path Shadowing Monte-Carlo*, illustrated in Fig. 1, is a Monte-Carlo method on shadowing paths. It is a predictive method since shadowing paths span over the future. Unlike a standard Monte-Carlo method, not all paths should have the same weight since  $\|h(x_{\text{past}}^k) - h(\tilde{x}_{\text{past}})\|$  is not uniform in  $k$ . This is to say, certain paths *shadow* more accurately  $x_{\text{past}}$  than others and should be considered as more *likely* to be extensions of the observed  $\tilde{x}_{\text{past}}$ . Path Shadowing Monte-Carlo estimates (10) by averaging  $q(x)$  on paths  $x^1, \dots, x^n$  with weights  $w_1, \dots, w_n$ . This yields the following estimator

$$\frac{1}{n} \sum_{k=1}^n w_k q(x^k), \quad (12)$$

called the Nadaraya–Watson estimator [33], [34]. In the following, we choose Gaussian weights, to wit

$$w_k = c \exp \left[ -\frac{\|h(x_{\text{past}}^k) - h(\tilde{x}_{\text{past}})\|^2}{2\eta^2} \right]$$

with  $c$  such that  $\frac{1}{n} \sum_{k=1}^n w_k = 1$ . The set of shadowing paths  $H_\eta(\tilde{x}_{\text{past}})$  (11) can be defined as the set of all paths that are one standard deviation away from  $\tilde{x}_{\text{past}}$  for the Gaussian kernel.

The following theorem proves the convergence of the estimator (12) under standard hypotheses.

### Theorem 1 (Path Shadowing Monte-Carlo Method)

*If  $q$  is continuous with  $\mathbb{E}\{|q(x)|\} < +\infty$  and the distribution  $p$  of  $x$  is continuous with  $p(x) > 0$  for all  $x \in \mathbb{R}^N$ , then given  $x^1, \dots, x^n, \dots$  independent realizations of  $x$ , the Path Shadowing Monte-Carlo estimator with  $h(x) = x$  converges almost surely*

$$\frac{1}{n} \sum_{k=1}^n w_k q(x^k) \longrightarrow \mathbb{E}\{q(x) \mid x_{\text{past}} = \tilde{x}_{\text{past}}\}$$

for a certain limit  $n \rightarrow +\infty$  and  $\eta \rightarrow 0$ .

The proof is in Appendix E. The continuity assumptions as well as the assumption that  $p > 0$  are technical and can be softened; note that a  $p_\theta$  of the form (1) satisfies them. We refer the reader to [35] for convergence theorems under more generic hypotheses.

Path Shadowing Monte-Carlo can be seen as a kernel method on log-price paths using a Gaussian kernel. It is a fully non-local method in the sense that the collected paths  $x^k$  may be far away in the past from  $\tilde{x}_{\text{past}}$ , possibly in a generated dataset of paths, contrary to non-local means [36] that only consider neighborhoods of a patch. Such non-locality ensures in practice the independence condition in theorem 1.

### B. Generating Shadowing Paths

Collecting enough shadowing paths from the historical realization of S&P is unrealistic. Indeed, financial price paths are noisy, two different observed paths are likely to be distant from one another, signifying a high-entropy underlying process. Thus, on limited data, the set  $H_\eta(\tilde{x}_{\text{past}})$  will contain almost no paths for reasonable values of  $\eta$ , required to be small for the method to converge.

We thus propose to scan for paths  $x^1, \dots, x^n$  in a long generated dataset of log-prices, allowing us to take  $n \gg 1$  and selective shadowing threshold  $\eta \ll 1$ . This however immediately introduces a modelling error, due to the fact that we are estimating (10) where  $\mathbb{E}$  is now the expectation with regard to the model distribution and not the true distribution  $p(x)$ .

A good model should generate trajectories that capture the long-range dependencies in order for the shadowing paths to have predictive power on the future of  $\tilde{x}_{\text{past}}$ . It should also generate a variety of realistic scenarios in order to find enough paths in  $H_\eta(\tilde{x}_{\text{past}})$  for small  $\eta$ . As discussed in section II the SS model precisely meets these requirements: it promotes diversity of generated paths through entropy maximization while reproducing main statistical properties of financial prices. Furthermore, should our generative algorithm produce occasionally unrealistic paths, they would be given a very small weight  $w$  and would not contribute to the estimation of  $q(x)$ . Shadowing Monte-Carlo is thus robust to outliers in the generated set of paths. Note also that this dataset can be

generated once and can be scanned again and again for several prediction dates.

A crucial point for PS-MC to give good results is to understand how the path embedding  $h$  affects the notion of path proximity. Such embedding should be chosen adequately. It should pick relevant features of  $x_{\text{past}}$  to predict the quantity of interest  $q(x)$ , while remaining low-dimensional such that enough paths can be found in  $H_\eta(\tilde{x}_{\text{past}})$  for small  $\eta$ . The naive embedding  $h(x) = (\delta x(t), -M + 1 \leq t \leq 0) \in \mathbb{R}^M$  limits the past horizon  $M$ , since  $M$  is then also the dimensionality of  $h$ .

We propose an embedding  $h$  that again leverages the scale-invariance of  $x$  in the same way our Scattering Spectra framework does, and incorporates the influence of distant past while remaining low-dimensional. It consists of multi-scale increments in the past with a power-law decaying weight

$$h_{\alpha,\beta}(x) = \left( \frac{x(t) - x(t-\ell)}{\ell^\beta}, \ell = \lfloor \alpha^m \rfloor, m = 1, 2, \dots \right) \quad (13)$$

for a certain  $\alpha > 1$  so that the past is progressively coarse-grained, and  $\beta \geq 0$  so that the far away past is progressively discounted. For a given  $\tilde{x}_{\text{past}}$  we choose  $\eta$  to be equal to  $\hat{\eta} \|h(\tilde{x}_{\text{past}})\|$ , which amounts to normalize the distance  $\|h(x_{\text{past}}) - h(\tilde{x}_{\text{past}})\|$  by  $\|h(\tilde{x}_{\text{past}})\|$ . Such  $h$  is a discretization of a continuous embedding that satisfies scaling and dilation equivariance, see Appendix F. In practice we truncate the progression  $\lfloor \alpha^1 \rfloor, \lfloor \alpha^2 \rfloor, \dots$  in order for the span of  $h$  to be bounded by 126 trading days in the past (corresponding to half a year).

The main parameters are thus  $\alpha, \beta$  and  $\hat{\eta}$ . Parameter  $\alpha$  determines the dimensionality of the path embedding  $h_{\alpha,\beta}$ , small  $\alpha$  yields high-dimensional embedding. Parameter  $\beta$  (not to be confused with the parameters  $\beta_{0,1,2}$  of the PDV model) rules the relative importance of distant past to recent past in the selection of shadowing paths. Large  $\beta > 0$  means that the recent past bears more weight. Finally, there is a bias-variance trade-off in the choice of  $\hat{\eta}$ . When  $\hat{\eta} \ll 1$  only the closest path will be used for averaging, leading to large variance estimator (12). When  $\hat{\eta} \gg 1$  then all paths are averaged uniformly, including paths whose past  $x_{\text{past}}$  has nothing to do with  $\tilde{x}_{\text{past}}$ , thus deteriorating the bias of our PS-MC estimator.

### C. Volatility Prediction

As a meaningful application of Path Shadowing Monte-Carlo, we consider in this section the prediction of the future daily realized variance over  $T$  days, for several values of  $T$ :

$$q_T(x) = \frac{252}{T} \sum_{u=t+1}^{t+T} |\delta x_u|^2.$$

We consider all 2112 dates  $t$  from January 2015 to March 2023. For each of them we consider the realized variance over  $T = 1, 7, 25, 75, 150$  days.

Our PS-MC method (12) uses paths generated from the model presented in section II. We compute the Scattering Spectra statistics (9) on observed 3772 S&P log-prices from January 2000 to December 2014, such that all our predictions are *out-of-sample*. From such statistics we generate 32 768

trajectories of same size 3772 (see Fig. 8 for examples), that represents  $n \approx 115$  million paths  $x^1, \dots, x^n$  of size 126+150 days. For a given  $\tilde{x}_{\text{past}}$  we scan such dataset and select the 50 000 closest paths in the sense of the distance induced by embedding (13), parameterized by  $\alpha, \beta$ . While this scanning step is computationally demanding, it can be fully parallelized. We then perform a weighted average on those closest paths, parameterized by  $\hat{\eta}$ .

Parameters  $\alpha, \beta, \hat{\eta}$  are calibrated using our generative model itself, in order to avoid any over-fitting on the limited train data from S&P. We choose those parameters such that  $q_T(x)$  is optimally predicted within the model. More precisely, we take 1100 snippets  $\tilde{x}_{\text{past}}$  from the generated dataset, for which we have access to  $\tilde{x}_{\text{future}}$ . We obtain an estimate of  $q_T$  for these 1100 dates and  $T = 7, 25, 75, 150$ . We then choose the best  $\alpha, \beta, \hat{\eta}$  so as to maximize the  $R^2$  score between prediction and true values of the SS model itself. This yields the following optimal parameters:  $\alpha = 1.15, \beta = 0.9, \hat{\eta} = 0.075$  and a path embedding  $h_{\alpha,\beta}$  of dimension 34.

Let us note that these optimal values barely change when predicting realized variance at different maturities. This is because of the scale-invariance of both the model and of the path embedding. Note also that  $\alpha = 1.15$  in (13) means that the values  $\ell = 1, 2, 3, 4$  appear multiple times. This amounts to ascribing an even larger weight to small time lags.

Number of days $T$	1	7	25	75	150
Benchmark	-0.16	0.43	-0.05	-0.58	-0.79
PDV (SS regression)	0.35	0.53	0.17	-0.44	-0.95
PDV (linear regression)	<b>0.36</b>	0.55	0.29	0.0	-0.07
SS Path Shadowing	0.32	<b>0.56</b>	<b>0.33</b>	<b>0.07</b>	<b>0.01</b>

TABLE I: Prediction of realized daily volatility ( $R^2$  score). The PS-MC method based on the Scattering Spectra (SS) model outperforms the recently introduced PDV model at all time-scales  $\geq 7$  days, and upholds predictive power up to a period of  $\approx 150$  days. For  $T = 1$  day, however, the PDV model performs best. The benchmark is the realized variance on the  $T$  previous days. The PDV model was calibrated using two different methods, see Appendix D.

The prediction quality is measured through the  $R^2$  score on future volatility estimates and are shown in Table I for different maturities  $T$ . As a simple benchmark we consider the realized variance on the  $T$  previous days as a predictor of  $q_T(x)$ . As a second, more challenging, benchmark we consider the recent Path-Dependent Volatility (PDV) model introduced in [9], which reads

$$\sqrt{q_T(x)} = \beta_0 + \beta_1 R_{1,t} + \beta_2 \sqrt{R_{2,t}}$$

$$\text{with } R_{1,t} = k_1 \star \delta x(t), \quad R_{2,t} = k_2 \star |\delta x|^2(t),$$

where  $k_1$  and  $k_2$  are both linear combinations of exponential kernels, acting on past returns and past absolute returns<sup>3</sup>. We take the very same kernels as specified in [9] but we determine

<sup>3</sup>Note that the authors of Ref. [9] estimate realized daily variance using 5-min ticks for better estimation, but the scores we obtain with daily ticks are actually similar for the longest maturities  $T = 3$  and  $T = 5$  that were tested in their study. Hence, we do not think that using 5-min ticks would substantially change the conclusions reached below for  $T \geq 7$  days.



the linear regression coefficients  $\beta_0, \beta_1, \beta_2$  in two different ways, see Appendix D. The first one amounts to reproducing as well as possible the Scattering Spectra, as we did in section III to reproduce the average smile. The corresponding values of  $\beta_{0,1,2}$  are given in Table IIIa. Unfortunately, this leads to rather poor results as far as volatility prediction is concerned, specially for large  $T$ , see Table I.

In order to favour the PDV model as much as possible, we used a different calibration aiming at best regressing  $\sqrt{qT}$  for each maturity  $T$  separately, on the same train period as for the SS model, i.e. from January 2000 to December 2014, see Table IIIb for calibrated values.

Using the very same shadowing paths for all maturities, our Path Shadowing Monte-Carlo method outperforms both the naive benchmark and the PDV model for all maturities from  $T > 7$  days, and ties with PDV for  $T = 7$  days, see Table I. In particular, our method upholds predictive power up to  $\approx 150$  days, which none of the two other methods are capable of. This is, we believe, due to the scale-invariance of both the SS model and our choice of path embedding (13). Again, we insist on the fact that the PDV model parameters are refitted for each maturity  $T$  whereas the SS model is calibrated once and for all. Of course, the SS model contains many more parameters than the PDV model, and we have not attempted to modify the shape of the kernel  $k_1$  and  $k_2$  introduced in [9], so there is clearly room for improvements there as well, which could be investigated in future work.

These results vindicate both the generative model presented in section II and the PS-MC method of the present section. In particular the Scattering Spectra generative model, now based on 182 parameters determined on the period 2000-2014, is *not* over-fitting the training dataset.

## V. OPTION PRICING & TRADING GAMES

In section III, we have used the Hedged Monte-Carlo method to price options *unconditionally* within the SS model, i.e. by averaging over all possible price paths of a given length  $T$ . This allowed us to obtain *average smiles* as a function of maturity, which we compared to those obtained using the same procedure but with real S&P trajectories.

Now, at a given date, option prices reflect anticipations of the market, conditioned on present market conditions – in particular the past price trajectory – and any available information about the future, such as earning announcements, dividends, political events, etc. Of course, such events cannot be directly captured by a purely statistical model, however faithful it might be. Still, it is interesting to run the exercise of pricing option smiles that anticipates the future solely based on the past of underlying price process.

In this section we investigate this question by combining the SS model presented in section II with Path Shadowing introduced in section IV. Option prices are then obtained by upgrading the Hedged Monte-Carlo method [29] with, as an input, shadowing paths generated by the model. The overall level of the resulting smiles at time  $t$  is nothing but the prediction of the future realized volatility for  $t' \in [t, t + T]$ , which was already investigated in the previous section. We

now extend such prediction to the full shape of the smile. We assess the quality of our smiles by trading a buy-sell signal on options whenever the model option smile is telling us that the option is under-priced or over-priced compared to another smile model, or of the option market itself.

### A. Path Shadowing Hedged Monte-Carlo

Hedged Monte-Carlo (HMC) [29] enables one to use time series of prices to compute the option prices. It iteratively determines the optimal price and hedging strategy by minimizing the expected financial risk of a portfolio containing the option to be priced and its hedge, at all times  $t = T - 1, T - 2, \dots, 0$ . The expected risk is computed as an average over paths, which in the present context are the shadowing paths, based on the notion of distance induced by the path embedding  $h$  (13). This defines the Path Shadowing Hedged Monte-Carlo (PS-HMC) that can be used in a versatile way to price any derivative contract. We use the same Gaussian weights (12) and the very same parameters  $\alpha, \beta, \eta$  detailed in section IV-C, that are optimal for volatility prediction within the model itself.

Fig. 4 shows the resulting smiles as a function of rescaled log-moneyness, for different maturities and at 3 typical dates. As one would have hoped, the level, but also the slope and the curvature of those smiles do depend on the chosen date, and more precisely on the actual path trajectory of the price before that date.




Fig. 4: Conditional smiles in a Scattering Spectra (SS) model at 3 different dates 2018-10-23, 2019-06-14 and 2020-02-26. They are computed through hedged Monte-Carlo on shadowing paths generated using the SS model. Note that the level, the slope and the curvature of those smiles all strongly depend on the chosen date.

### B. A Trading Game

In order to assess the quality of the smiles predicted by the SS model, we set up the following trading game. We trade call options at several dates  $t$  on the option market. We neglect bid-ask spread and consider the option price  $\mathcal{C}^{\text{MKT}}(t, T, K)$  to be the quoted mid-price. We denote  $\sigma^{\text{MKT}}(t, T, K)$  the observed implied volatility computed from  $\mathcal{C}^{\text{MKT}}(t, T, K)$  and  $\sigma^{\text{model}}(t, T, K)$  the implied volatility computed within the model that we decide to trade with.

We test the following trading strategy: buy the corresponding option from the market whenever we deemed it under-priced, i.e.  $\sigma^{\text{MKT}}(t, T, K) < \sigma^{\text{model}}(t, T, K)$  or sell it if we

deem it over-priced, i.e.  $\sigma^{\text{MKT}}(t, T, K) > \sigma^{\text{model}}(t, T, K)$ . We then follow the hedged option until maturity and register the corresponding profit or loss associated to the trade.

The buy-sell signal of such strategy is thus given by

$$\epsilon_t = \begin{cases} +1 & \text{if } \sigma^{\text{MKT}}(t, T, K) < \sigma^{\text{model}}(t, T, K), \\ -1 & \text{if } \sigma^{\text{MKT}}(t, T, K) > \sigma^{\text{model}}(t, T, K). \end{cases}$$

The un-hedged forward P&L $_t(T, K)$  of one transaction is obtained as

$$\text{P\&L}_t(T, K) = v_t \epsilon_t \left( (e^{x_t+T} - K)_+ - C^{\text{MKT}}(t, T, K) \right)$$

where  $v_t$  is the volume of option traded. In order to remove non-stationary effect due to the long-term change in the value of the underlying, we take  $v_t = 100/S_t = 100e^{-x_t}$  which amounts to trade options on percentage of variation of the underlying rather than the underlying itself.

To reduce the variance of the strategy, we hedge the option using a simple Black-Scholes delta-hedge with a constant volatility 0.2. Such delta-hedge gives zero profit on average but reduces greatly (though not optimally, see [29], [37] for details) the variance of P&L $_t$ .

In the following, we will consider the model smile  $\sigma^{\text{model}}$  to be given either by the smile computed in the SS model using PS-HMC (model=SS) or the smile computed in a Path-Dependent Volatility model [9] (model=PDV), both using HMC. The two models (PDV and SS) are calibrated using the same data in the train period i.e. January 2000-December 2014. As in the previous section, the PDV model parameters are furthermore *re-optimized* for each maturity  $T$  independently as in Table IIIb, whereas the SS model is parameterized once and for all with the Scattering Spectra determined in the train period.

The trading game is then in both cases played out-of-sample, for all 2112 dates  $t$  from January 2015 to May 2023. We choose 5 maturities  $T = 8, 25, 50, 75, 150$  and 9 strikes at constant rescaled log-moneyness  $\mathcal{M} = -2, -1.5, -1, -0.5, 0, 0.5, 1, 1.5, 2$ .

Detailed P&Ls over 3 different periods of 3 year each are shown in Fig. 5. Their variance across dates  $t$  is shown in Fig. 12 in Appendix G. For most maturities and strikes, the trading game using the SS model yields positive P&Ls and clearly outperforms the trading game using the PDV model. In fact, one can directly play the SS model against the PDV model without any reference to the actual option market, fully confirming that the SS model outperforms PDV model for almost all maturities and strikes, see Appendix G and in particular Fig. 14.

Since the P&Ls are of the same order of magnitude across different strikes and maturities, we average them over all the maturities and strikes. Table II shows such grand averages and reveals that the trading game using the SS model is significantly more profitable than using the PDV model, with furthermore less variance across different periods. This is confirmed by the aggregated P&Ls over time, see Fig. 6.

	full period	2015-2017	2018-2020	2021-2023
PDV	$0.071 \pm 0.03$	$0.13 \pm 0.01$	$0.0 \pm 0.04$	$0.10 \pm 0.02$
SS	<b><math>0.13 \pm 0.03</math></b>	<b><math>0.14 \pm 0.02</math></b>	<b><math>0.14 \pm 0.04</math></b>	<b><math>0.11 \pm 0.02</math></b>

TABLE II: Average P&L of the trading game against the S&P market using the PDV model (re-optimized for each  $T$ ) or using the Scattering Spectra (SS) model, both calibrated using data from 2000 to 2014. Detailed P&Ls are shown in Fig. 5.




Fig. 5: P&Ls of the trading game against the S&P market with a PDV model (re-optimized for each  $T$ ) or with the Scattering Spectra (SS) model, both calibrated using data from 2000 to 2014. Each heatmap corresponds to a 3 years period, from top to bottom (2015-2017, 2018-2020, 2021-2023).

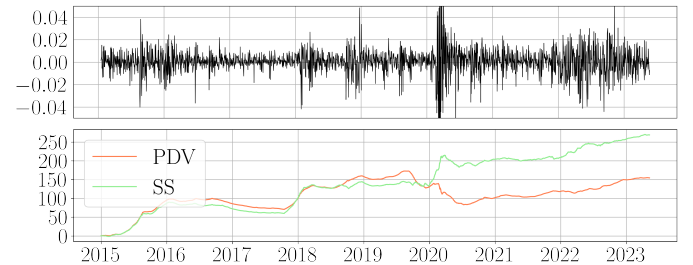


Fig. 6: (Bottom) Aggregated P&Ls of the trading game against the S&P market, played from 2015 to 2023, with a PDV model (re-optimized for each  $T$ ) or with the Scattering Spectra (SS) model, both calibrated using data from 2000 to 2014. (Top) Time series of S&P log-return on the trading game period.

## VI. CONCLUSION

We presented a statistical model of financial prices based on the Scattering Spectra introduced in our previous work [1]. Scattering Spectra are multi-scale analogues of the standard skewness and kurtosis. Such a model achieves a tradeoff between accuracy and diversity. It captures main statistical properties of prices as well as several scaling properties of option smiles. As a maximum entropy model with a small number of constraints, our Scattering Spectra model also avoids over-fitting.

We then introduced Path Shadowing Monte-Carlo (PS-MC) which enables building models of forward looking probabili-

ties  $p(x|x_{\text{past}})$  from any generative model of  $p(x)$ . Combined with our statistical model of prices, PS-MC provides state-of-the-art volatility predictions. Shadowing paths can also be used to obtain option smiles that depend only on the distribution of the price process. A trading game allowed us to show that the Scattering Spectra model correctly anticipates future price movements and favourably competes with the (simplest version of) the recently introduced Path-Dependent Volatility model. Extending our analysis to other underlying assets, such as single stocks, is of course important to validate our conclusions when price moves are more idiosyncratic.

Beyond prediction, Path Shadowing may be a way of tackling other burning questions in Finance, such as *typicality*: how typical or atypical is a given sequence of price changes? As far as Path Shadowing Monte-Carlo is concerned, one limitation of the method is that it requires to scan a large dataset of generated paths for delivering good performances. This scanning step could be done more efficiently. In particular, could one find ‘typical’ price paths that should be frequently selected for prediction in order to save intensive scanning efforts? Another highly relevant extension is towards the description of multivariate time series. We hope to address these issues in the near future using the methods introduced in this work.

*Acknowledgments:* We want to thank C. Aubrun, N. Cuvelle-Magar, S. Crépey, J. Gatheral, F. Guth, J. Guyon, B. Horvath, G. Pagès, M. Potters & V. Vargas for many insightful conversations on these topics.

## REFERENCES

- [1] R. Morel, G. Rochette, R. Leonarduzzi, J.-P. Bouchaud, and S. Mallat, “Scale dependencies and self-similar models with wavelet scattering spectra,” *arXiv preprint arXiv:2204.10177*, 2022.
- [2] A. v. d. Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, “Wavenet: A generative model for raw audio,” *arXiv preprint arXiv:1609.03499*, 2016.
- [3] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, “Attention is all you need,” *Advances in neural information processing systems*, vol. 30, 2017.
- [4] Q. Wen, T. Zhou, C. Zhang, W. Chen, Z. Ma, J. Yan, and L. Sun, “Transformers in time series: A survey,” *arXiv preprint arXiv:2202.07125*, 2022.
- [5] S. Heston, “A closed-form solution for options with stochastic volatility with applications to bond and currency options,” *Review of Financial Studies*, vol. 6, pp. 327–343, 1993.
- [6] E. Bacry, A. Kozhemyak, and J.-F. Muzy, “Log-normal continuous cascade model of asset returns: aggregation properties and estimation,” *Quantitative Finance*, vol. 13, no. 5, pp. 795–818, 2013.
- [7] J. Gatheral, T. Jaisson, and M. Rosenbaum, “Volatility is rough,” *Quantitative finance*, vol. 18, no. 6, pp. 933–949, 2018.
- [8] P. Wu, J.-F. Muzy, and E. Bacry, “From rough to multifractal volatility: The log s-fbm model,” *Physica A: Statistical Mechanics and its Applications*, vol. 604, p. 127919, 2022. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0378437122005866>
- [9] J. Guyon and J. Lekeufack, “Volatility is (mostly) path-dependent,” *Quantitative Finance*, pp. 1–38, 2021.
- [10] M. Fukasawa, “Short-time at-the-money skew and rough fractional volatility,” *Quantitative Finance*, vol. 17, no. 2, pp. 189–198, 2017.
- [11] M. Potters, R. Cont, and J.-P. Bouchaud, “Financial markets as adaptive systems,” *Europhysics letters*, vol. 41, no. 3, p. 239, 1998.
- [12] D. V. Anosov, *Geodesic flows on closed Riemann manifolds with negative curvature*. American Mathematical Society, 1969, no. 89-90.
- [13] Y. G. Sinai, “Gibbs measures in ergodic theory,” *Russian Mathematical Surveys*, vol. 27, no. 4, p. 21, 1972.
- [14] R. Bowen, “ $\omega$ -limit sets for axiom a diffeomorphisms,” *Journal of differential equations*, vol. 18, no. 2, pp. 333–339, 1975.
- [15] S. M. Hammel, J. A. Yorke, and C. Grebogi, “Do numerical orbits of chaotic dynamical processes represent true orbits?” *Journal of Complexity*, vol. 3, no. 2, pp. 136–145, 1987.
- [16] C. Salvi, M. Lemerrier, C. Liu, B. Horvath, T. Damoulas, and T. Lyons, “Higher order kernel mean embeddings to capture filtrations of stochastic processes,” *Advances in Neural Information Processing Systems*, vol. 34, pp. 16635–16647, 2021.
- [17] A. Alden, C. Ventre, B. Horvath, and G. Lee, “Model-agnostic pricing of exotic derivatives using signatures,” in *Proceedings of the Third ACM International Conference on AI in Finance*, 2022, pp. 96–104.
- [18] J. Delemotte, S. D. Marco, and F. Segonne, “Yet another analysis of the sp500 at-the-money skew: crossover of different power-law behaviours,” *Available at SSRN 1520443*, 2023.
- [19] B. B. Mandelbrot, A. J. Fisher, and L. E. Calvet, “A multifractal model of asset returns,” 1997.
- [20] L. Borland, J.-P. Bouchaud, J.-F. Muzy, and G. Zumbach, “The dynamics of financial markets—mandelbrot’s multifractal cascades, and beyond,” *arXiv preprint cond-mat/0501292*, 2005.
- [21] L. Bergomi, “Smile dynamics iv,” *Available at SSRN 1520443*, 2009.
- [22] V. Vargas, T.-L. Dao, and J.-P. Bouchaud, “Skew and implied leverage effect: smile dynamics revisited,” *International Journal of Theoretical and Applied Finance*, vol. 18, no. 04, p. 1550022, 2015.
- [23] E. T. Jaynes, “Information Theory and Statistical Mechanics,” *Physical Review*, vol. 106, no. 4, pp. 620–630, May 1957.
- [24] R. Lustig, “Microcanonical Monte Carlo simulation of thermodynamic properties,” *The Journal of Chemical Physics*, vol. 109, no. 20, pp. 8816–8828, 11 1998. [Online]. Available: <https://doi.org/10.1063/1.477552>
- [25] M. Betancourt, “A conceptual introduction to hamiltonian monte carlo,” *arXiv preprint arXiv:1701.02434*, 2017.
- [26] G. Battle, “A block spin construction of ondelettes. part i: Lemarié functions,” *Communications in Mathematical Physics*, vol. 110, no. 4, pp. 601–615, 1987.
- [27] P.-G. Lemarié, “Ondelettes à localisation exponentielle,” *J. Math. Pures Appl.*, vol. 67, pp. 227–236, 1988.
- [28] J.-P. Bouchaud, L. De Leo, V. Vargas, and S. Ciliberti, “Smile in the low moments,” *Risk*, vol. 26, no. 7, p. 56, 2013. [Online]. Available: <https://arxiv.org/pdf/1203.5703.pdf>
- [29] M. Potters, J.-P. Bouchaud, and D. Sestovic, “Hedged monte-carlo: low variance derivative pricing with objective probabilities,” *Physica A: Statistical Mechanics and its Applications*, vol. 289, no. 3, pp. 517–525, 2001. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0378437100005549>
- [30] M. Fukasawa, “Volatility has to be rough,” *Quantitative Finance*, vol. 21, no. 1, pp. 1–8, 2021.
- [31] C. Bayer, P. Friz, and J. Gatheral, “Pricing under rough volatility,” *Quantitative Finance*, vol. 16, no. 6, pp. 887–904, 2016.
- [32] J. Guyon and M. El Amrani, “Does the term-structure of equity at-the-money skew really follow a power law?” *Available at SSRN 4174538*, 2022.
- [33] E. A. Nadaraya, “On estimating regression,” *Theory of Probability & Its Applications*, vol. 9, no. 1, pp. 141–142, 1964.
- [34] G. S. Watson, “Smooth regression analysis,” *Sankhyā: The Indian Journal of Statistics, Series A*, pp. 359–372, 1964.
- [35] B. E. Hansen, “Uniform convergence rates for kernel estimation with dependent data,” *Econometric Theory*, vol. 24, no. 3, pp. 726–748, 2008.
- [36] A. Buades, B. Coll, and J.-M. Morel, “Non-local means denoising,” *Image Processing On Line*, vol. 1, pp. 208–212, 2011.
- [37] J.-P. Bouchaud and M. Potters, *Theory of financial risk and derivative pricing: from statistical physics to risk management*. Cambridge university press, 2003.
- [38] J. Bruna and S. Mallat, “Multiscale sparse microcanonical models,” *Mathematical Statistics and Learning*, vol. 1, no. 3, pp. 257–315, 2019.
- [39] L. C. Evans, *Partial differential equations*. American Mathematical Society, 2022, vol. 19.

## APPENDIX A SAMPLING ALGORITHM

Drawing samples from the Scattering Spectra model (1) could be performed through an exact algorithm [24]. This is however computationally expensive when the number of statistics  $\Phi$  is not small. Instead we use an approximate

algorithm based on gradient descent [38]. We consider the set of trajectories  $x$  having approximately the correct statistics

$$\Omega_\epsilon = \{x \in \mathbb{R}^N \mid \|\Phi(x) - \Phi(\tilde{x})\| \leq \epsilon\}.$$

The width  $\epsilon$  accounts for the statistical error between  $\Phi(\tilde{x})$  and the true  $\mathbb{E}\{\Phi(x)\}$ . We chose  $\epsilon = 10^{-3}\|\Phi(\tilde{x})\|$ . An approximate sampling of model (1) is performed by sampling trajectories in  $\Omega_\epsilon$  [38]. The algorithm is initialized with a realization  $x \in \mathbb{R}^N$  of a Gaussian noise which has a maximum entropy distribution. Then we minimize the loss  $\ell(x) = \|\Phi(x) - \Phi(\tilde{x})\|^2$  through gradient descent

$$x \leftarrow x - \gamma \nabla \ell(x).$$

until the condition  $\ell \leq \epsilon$  is reached. The gradient descent is implemented with the L-BFGS-B algorithm.

## APPENDIX B WAVELETS

We impose that the wavelet  $\psi$  satisfies the following energy conservation law called the Littlewood-Paley equality

$$\forall \omega > 0, \quad \sum_{j=-\infty}^{+\infty} |\widehat{\psi}(2^j \omega)|^2 = 1. \quad (14)$$

A Battle-Lemarié wavelet [26], [27] is an example of such wavelet, shown in Fig. 7. It has an exponential decay away from  $t = 0$  and has  $m = 4$  vanishing moments




Fig. 7: Left: complex Battle-Lemarié wavelet  $\psi(t)$  as a function of  $t$ . Right: Fourier transform  $\widehat{\psi}(\omega)$  as a function of  $\omega$ .

The wavelet transform is computed up to a largest scale  $2^J$  which is smaller than the signal size  $N$ . The signal lower frequencies in  $[-2^{-J}\pi, 2^{-J}\pi]$  are captured by a low-pass filter  $\varphi_J(t)$  whose Fourier transform is

$$\widehat{\varphi}_J(\omega) = \left( \sum_{j=J+1}^{+\infty} |\widehat{\psi}(2^j \omega)|^2 \right)^{1/2}. \quad (15)$$

One can verify that it has a unit integral  $\int \varphi_J(t) dt = 1$ . To simplify notations, we write this low-pass filter as a last scale wavelet  $\psi_{J+1} = \varphi_J$ , and  $W_{J+1}x(t) = x \star \psi_{J+1}(t)$ . By applying the Parseval formula, we derive from (14) that for all  $x$  with  $\|x\|^2 = \int |x(t)|^2 dt < \infty$

$$\|Wx\|^2 = \sum_{j=-\infty}^{J+1} \|x \star \psi_j\|^2 = \|x\|^2.$$

The wavelet transform  $W$  preserves the norm and is therefore invertible, with a stable inverse.

Properties of signal increments are carried over to wavelet coefficients by observing that wavelet coefficients are obtained by filtering signal increments  $\delta_j x(t) = x(t) - x(t - 2^j)$  with a dilated integrable filter:

$$x \star \psi_j(t) = \delta_j x \star \theta_j(t) \quad \text{where} \quad \theta_j(t) = 2^{-j} \theta(2^{-j}t), \quad (16)$$

where filter  $\theta$  is obtained from  $\psi$  through  $\widehat{\theta}(\omega) = \widehat{\psi}(\omega) / (1 - e^{-i\omega})$ . This is because  $1 - e^{-i2^j \omega}$  is the Fourier transform of  $\delta(t) - \delta(t - 2^j)$ , the filter that creates increments. From (16) we get that if  $\delta_j x(t)$  is stationary then  $x \star \psi_j(t)$  is also stationary.




Fig. 8: (Top) Historical S&P log-returns from 2000 to 2014. (Bottom) Examples of generated syntheses from the Scattering Spectra model, which are scanned for shadowing paths.

## APPENDIX C

### SMILE SENSITIVITY IN THE SCATTERING SPECTRA MODEL

The Scattering Spectra model defined in (1) depends only on the estimated values  $\Phi(\tilde{x})$  of the Scattering Spectra estimated on a single realization  $\tilde{x}$  of S&P. This models gives an unconditional smile shown in Fig. 3b.

We are interested in the change in this unconditional smile in the case where the statistics  $\Phi(\tilde{x})$  change significantly. Thanks to the interpretation of Scattering Spectra coefficients, we can see what happens to the smile if the market is ‘more skewed’ or ‘more kurtic’ for example.

Here we focus on the shape of the smiles. Of course, changing the amplitude of  $\Phi_2$ , which is equivalent to increasing the overall volatility of the model, only moves the overall level of the smile up or down.

These sensitivities in general can be seen as amplifying or reducing the departure of the price process from a Gaussian process  $x_{\text{Gaussian}}$ , also called Black-Scholes model, with the same average volatility, meaning the same value of  $\Phi_2(\tilde{x})$ .

$$\text{new statistics} = (1 - \lambda)\Phi(x_{\text{Gaussian}}) + \lambda\Phi(\tilde{x}).$$

If  $\lambda < 1$  the corresponding model should be ‘closer’ to a Black-Scholes model, if  $\lambda > 1$  the corresponding model should become less ‘Gaussian’. Fig. 9 shows 4 directions of change that are detailed below.

**Skewness**  $|\Phi_3(x)|$ . The skewness coefficients  $\Phi_3(x_{\text{Gaussian}})$  should be zero for a Gaussian process. We consider a model of  $x$  with modified statistics

$$\Phi_3(x) = \lambda \Phi_3(\tilde{x})$$

for 3 values of  $\lambda$ . For  $\lambda = 0$ , the modeled process is not skewed, meaning that an increment trajectory  $\delta x$  is equally likely as a trajectory  $-\delta x$ . Unsurprisingly, we get smiles that are symmetrical at  $\mathcal{M} = 0$ . For  $\lambda = 1$  we get the same smile as in the Scattering Spectra model of S&P. For  $\lambda = 1.3$ , the smile has a higher downward slope, as expected.

**Time-asymmetry**  $\text{Im } \Phi(x)$ . Our representation  $\Phi$  is complex-valued. While  $\Phi_1, \Phi_2$  are real, skewness  $\Phi_3$  and kurtosis  $\Phi_4$  may have non-zero imaginary parts that were shown to characterize certain types of time-asymmetry. We consider a model  $x$  whose statistics are

$$\Phi(x) = \text{Re } \Phi(\tilde{x}).$$

It is thus time-reversible, i.e. a trajectory  $x(t)$  is equally likely as  $x(-t)$ . We notice that its smiles are symmetrical, but compared to the previous case, these are not symmetrical around  $\mathcal{M} = 0$  but around values  $\mathcal{M} > 0$  depending on the maturity. This is consistent since the process still has non-zero skewness  $|\Phi_3(x)|$ .

**Kurtosis**  $\Phi_1$ . For a Gaussian process,  $\Phi_1(x_{\text{Gaussian}}) = \pi/4$ . We consider a model  $x$  with modified statistics

$$\Phi_1(x) = (1 - \lambda) \frac{\pi}{4} + \lambda \Phi_1(\tilde{x}).$$

For  $\lambda = 0.5$ , the model is less kurtic than the S&P, it shows smiles that tend to flatten around a straight line with negative slope. For  $\lambda = 1.75$ , the model is more kurtic and the smiles have more curvature, as expected.

**Kurtosis**  $\Phi_4$ . We consider a model  $x$  with modified statistics

$$\Phi_4(x) = (1 - \lambda) \Phi_4(x_{\text{Gaussian}}) + \lambda \Phi_4(\tilde{x}).$$

The change in smiles for two different values  $\lambda = 0.5, 1.5$  seem small compared to the other effects presented, however such changes impacts a lot the trajectories.

## APPENDIX D

### THE PATH-DEPENDENT VOLATILITY MODEL

The Path-Dependent Volatility (PDV) model introduced in [9] consists of a 4-factor Markovian model with 9 parameters. Writing the price process as  $S_t = S_0 e^{x_t}$ , it assumes that

$$\begin{aligned} \frac{dS_t}{S_t} &= \sigma_t dW_t, \\ \sigma_t &= \sigma(R_{1,t}, R_{2,t}), \\ \sigma(R_1, R_2) &= \beta_0 + \beta_1 R_1 + \beta_2 \sqrt{R_2} \\ R_{1,t} &= \int_{-\infty}^t K_1(t-u) \frac{dS_u}{S_u} \\ R_{2,t} &= \int_{-\infty}^t K_2(t-u) \left( \frac{dS_u}{S_u} \right)^2 \end{aligned}$$

Among these parameters, 6 are used to parameterize the kernels  $K_1, K_2$  both being a linear combination of exponentials, and 3 are the regression coefficients  $\beta_0, \beta_1, \beta_2$ . The 6 kernel parameters are set to the optimal values calibrated on S&P and presented in [9] (Table 11).

To obtain unconditional smiles in the PDV model, the process is evolved with 1 step per day until it reaches a




Fig. 9: Smile sensitivity to change in Scattering Spectra statistics  $\Phi(\tilde{x})$ , decomposed as changes in skewness  $|\Phi_3|$ , time-asymmetry  $\text{Im } \Phi$ , kurtosis  $\Phi_1$  or kurtosis  $|\Phi_4|$ . A value  $\lambda < 1$  indicates respectively, no skewness, no time-asymmetry, less kurtosis. A factor  $\lambda = 1$  does not change the statistics  $\Phi(\tilde{x})$  estimated on S&P. Besides well-known influence of the skewness and kurtosis on the shape of the smile, the Scattering Spectra  $\Phi(x)$  also explains the contribution of time-asymmetry in the shape of the smile.

stationary regime. The 3 regression coefficients  $\beta_0, \beta_1, \beta_2$  are calibrated on the Scattering Spectra statistics  $\Phi$  and reported in Table IIIa. We choose the values that minimize the distance  $\|\Phi(x_{\text{PDV}}) - \Phi(\tilde{x})\|_2$  where  $\tilde{x}$  is the S&P log-price time series consisting of  $N = 5827$  days from January 2000 to April 2023 and  $x_{\text{PDV}}$  are 10 realizations of same size from the PDV model. The Scattering Spectra of the ‘closest’ PDV model to the S&P are shown on Fig.10.

	$\beta_0$	$\beta_1$	$\beta_2$
unconditional	0.050	-0.13	0.56

(a) Optimal parameters of the Path-Dependent Volatility model calibrated on Scattering Spectra statistics over the full period 2000-2023.

	$\beta_0(T)$	$\beta_1(T)$	$\beta_2(T)$
$T = 1$	0.022	-0.062	0.64
$T = 7$	0.035	-0.066	0.75
$T = 8$	0.036	-0.065	0.74
$T = 25$	0.051	-0.050	0.68
$T = 50$	0.066	-0.040	0.62
$T = 75$	0.079	-0.039	0.55
$T = 126$	0.098	-0.033	0.46
$T = 150$	0.10	-0.030	0.43

(b) Optimal parameters of the Path-Dependent Volatility model on the period 2000-2014 (train period), obtained through linear regression of future realized volatilities over  $T$  days.

TABLE III: Path-Dependent Volatility (PDV) model calibration.

It shows that the kurtosis  $\Phi_4(x)$  is significantly lower than for the S&P500 data, see also Fig. 3e. Looking at the log-return trajectories shown in Fig. 11 we indeed notice clear qualitative discrepancies: the PDV log-returns are less intermittent than the S&P ones. In Fig. 11 we see that

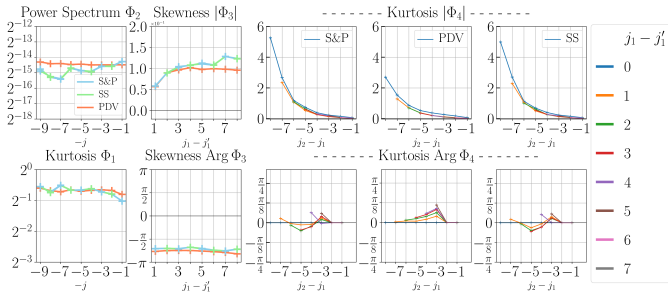


Fig. 10: Visualization of the Scattering Spectra statistics  $\Phi = (\Phi_1, \Phi_2, \Phi_3, \Phi_4)$ . These are estimated on the observed S&P500 realization, in a Path-Dependent Volatility (PDV) model and in the Scattering Spectra (SS) model.

structure functions (b) are poorly reproduced. Furthermore, the short term leverage effect is much too strong, see and, quite strikingly, Fig. 11c and Fig. 3d.

We could also choose parameters  $\beta_0, \beta_1, \beta_2$  as the one specified in [9], using optimal parameters for the regression of future realized volatility but it does not improve the results in Figs. 11,3. In the former case, the PDV trajectories feature a larger skewness  $\Phi_3$  than the S&P with abnormal negative log-return values, as shown in [9]. In the later case, the intermittency, measured in particular by  $\Phi_4$ , is much lower than the S&P.

This showcases the difficulty for a low-parametric model to capture a rich set of statistics such as the Scattering Spectra.

To obtain good volatility prediction and conditional smiles across different maturities  $T$ , to be used in trading games (see section V), we had to recalibrate the parameters  $\beta_0(T), \beta_1(T), \beta_2(T)$  for each maturity  $T$  independently, this in order to provide the best prediction of the future realized variance i.e. the overall level of the smile. Calibrated values are shown on Table IIIb.

#### APPENDIX E PROOF OF THEOREM 1

Let us write  $w_k = c_n g_\eta(x^k - \tilde{x})$  where

$$g_\eta(x) = (\eta\sqrt{2\pi})^{-N'} e^{-\frac{1}{2\eta^2}\|x_{\text{past}}\|^2}$$

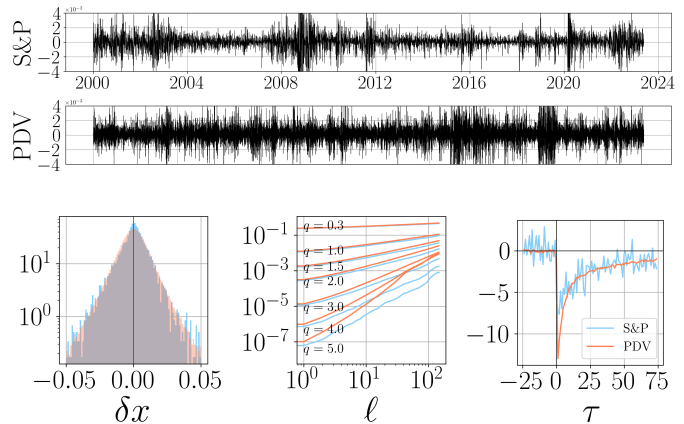
is a Gaussian kernel with  $N'$  being the dimension of  $x_{\text{past}}$  and  $c_n$  is such that  $\frac{1}{n} \sum_k w_k = 1$ . We write  $\hat{q}$  the estimator

$$\hat{q} = \frac{1}{n} \sum_{k=1}^n w_k q(x^k).$$

We prove the convergence of  $\hat{q}$  to  $\mathbb{E}\{q(x) \mid x_{\text{past}} = \tilde{x}_{\text{past}}\}$  almost surely by first taking the limit  $n \rightarrow +\infty$  and then  $\eta \rightarrow 0$ .

**Limit  $n \rightarrow \infty$ .** One can calculate  $c_n^{-1} = \frac{1}{n} \sum g_\eta(x^k)$  so that one has

$$\hat{q} = \frac{\frac{1}{n} \sum_{k=1}^n g_\eta(x^k - \tilde{x}) q(x^k)}{\frac{1}{n} \sum_{k=1}^n g_\eta(x^k - \tilde{x})}.$$



(a) Histogram (b) Struct. functions (c) Leverage

Fig. 11: S&P returns vs. synthetic returns using the Path-Dependent Volatility (PDV) model calibrated on Scattering Spectra (top), and standard statistics in the same PDV model (orange) of the S&P (blue): (a) Histogram of daily log-returns  $\delta x$ . (b) Structure functions  $\mathbb{E}\{|\delta_\ell x(t)|^q\}$ . (c) Leverage correlation  $\mathbb{E}\{\delta x(t - \tau)|\delta x(t)|^2\}$ . Compare with the results shown in Fig. 2.

Since  $g_\eta$  is bounded one has  $\mathbb{E}\{g_\eta(x - \tilde{x})|q(x)|\} < +\infty$  and  $\mathbb{E}\{g_\eta(x - \tilde{x})\} < +\infty$ . From the law of large numbers, knowing that  $\mathbb{E}\{g_\eta(x - \tilde{x})\} > 0$ , it follows

$$\hat{q} \xrightarrow{n \rightarrow +\infty} \frac{\mathbb{E}\{g_\eta(x - \tilde{x})q(x)\}}{\mathbb{E}\{g_\eta(x - \tilde{x})\}}.$$

**Limit  $\eta \rightarrow 0$ .** We will make use of the following lemma of approximation by convolution, proved for example in [39].

#### Lemma 2

If  $f \in \mathcal{C}^0 \cap L^1(dx)$  then for all  $\tilde{x} \in \mathbb{R}^N$

$$g_\eta \star f(\tilde{x}) \xrightarrow{\eta \rightarrow 0} \int f(\tilde{x}_{\text{past}}, x_{\text{future}}) dx_{\text{future}}.$$

Let us notice that  $\mathbb{E}\{g_\eta(x - \tilde{x})\} = \int g_\eta(\tilde{x} - x)p(x)dx = g_\eta \star p(\tilde{x})$ , and  $\mathbb{E}\{g_\eta(x - \tilde{x})q(x)\} = g_\eta \star (qp)(\tilde{x})$ . Since  $\mathbb{E}\{q(x)\} < +\infty$  one has  $qp \in L^1(dx)$ ,  $p$  being a probability distribution one also has  $p \in L^1(dx)$ . From the lemma we get:

$$\frac{\mathbb{E}\{g_\eta(x - \tilde{x})q(x)\}}{\mathbb{E}\{g_\eta(x - \tilde{x})\}} \xrightarrow{\eta \rightarrow 0} \frac{\int q(\tilde{x}_{\text{past}}, x_{\text{future}})p(\tilde{x}_{\text{past}}, x_{\text{future}})dx_{\text{future}}}{\int p(\tilde{x}_{\text{past}}, x_{\text{future}})dx_{\text{future}}},$$

where the denominator is non-zero because  $p(x) > 0$  for all  $x \in \mathbb{R}^N$ . The former term being  $\mathbb{E}\{q(x) \mid x_{\text{past}} = \tilde{x}_{\text{past}}\}$ , this proves the theorem.

#### APPENDIX F CHOICE OF PAST EMBEDDING $h$

The following proposition shows that the choice of  $h$  in our paper (13) induces equivariance properties on the set of shadowing paths  $H_\eta(\tilde{x}_{\text{past}})$ . We recall that we chose  $\eta = \hat{\eta}\|h(\tilde{x}_{\text{past}})\|$  for a fixed  $\hat{\eta}$ . These equivariance properties are proved on continuously sampled paths  $x_{\text{past}} = (x(t), t < 0)$

with  $t \in \mathbb{R}$ . In that case, we still write  $h_{\alpha,\beta}$  the continuously sampled analogue

$$h_{\alpha,\beta}(x) = \left( \frac{x(0) - x(-\ell)}{\ell^\beta}, \ell = \alpha^m, m \in \mathbb{Z} \right)$$

which is now of infinite dimension.

### Proposition 3

For  $h = h_{\alpha,\beta}$  with  $\alpha > 0$  and  $\beta > 1$  one has

1. (Multiplication equivariance) for  $\lambda > 0$

$$H_\eta(\lambda \tilde{x}_{\text{past}}) = \lambda \cdot H_\eta(\tilde{x}_{\text{past}})$$

2. (Dilation equivariance) writing  $\Gamma x(t) = x(\alpha t)$

$$H_\eta(\Gamma \tilde{x}_{\text{past}}) = \Gamma \cdot H_\eta(\tilde{x}_{\text{past}})$$

*Proof.* The first equivariance follows directly from the fact that  $h_{\alpha,\beta}$  is itself equivariant to multiplication  $h_{\alpha,\beta}(\lambda x_{\text{past}}) = \lambda h_{\alpha,\beta}(x_{\text{past}})$ . For dilation, one has

$$\begin{aligned} h_{\alpha,\beta}(\Gamma x_{\text{past}}) &= \left( \frac{x(0) - x(-\alpha\ell)}{\ell^\beta}, \ell = \alpha^m, m \in \mathbb{Z} \right) \\ &= \alpha^\beta \left( \frac{x(0) - x(-\ell)}{\ell^\beta}, \ell = \alpha^{m+1}, m \in \mathbb{Z} \right) \end{aligned}$$

This means that  $h_{\alpha,\beta}(\Gamma x_{\text{past}})$  is equal to  $h_{\alpha,\beta}(x_{\text{past}})$  up to a shift in indices and up to a multiplicative constant. It follows that

$$\|h_{\alpha,\beta}(\Gamma x_{\text{past}}) - h_{\alpha,\beta}(\Gamma \tilde{x}_{\text{past}})\| = \alpha^\beta \|h_{\alpha,\beta}(x_{\text{past}}) - h_{\alpha,\beta}(\tilde{x}_{\text{past}})\|.$$

Now, normalizing by  $\|h_{\alpha,\beta}(\Gamma \tilde{x}_{\text{past}})\| = \alpha^\beta \|h_{\alpha,\beta}(\tilde{x}_{\text{past}})\|$  yields  $H_\eta(\Gamma \tilde{x}_{\text{past}}) = \Gamma \cdot H_\eta(\tilde{x}_{\text{past}})$ .  $\square$

## APPENDIX G

### ADDITIONAL TRADING GAME STATISTICS

In addition to the P&Ls and aggregated P&Ls of a trading game played against the option market shown in Figs. 5,6, we show in Figs. 12,13 the standard deviation and winning rate, defined as the average number of times the payoff of a trade is positive.

In the remaining figures, we also show the statistical results of the trading game between PDV and SS model. These results do not require option market data (but require actual price series of the underlying) and directly test the relative quality of purely statistical price models. As seen in Fig. 14, trading game unequivocally favours the Scattering Spectra model framework over the Path-Dependent Volatility model.




Fig. 12: Standard deviation of P&Ls of a trading game playing the Scattering Spectra (SS) model vs. S&P or the Path-Dependent Volatility (PDV) model vs. S&P. Each heatmap corresponds to a 3 years period, from top to bottom (2015-2017, 2018-2020, 2021-2023).




Fig. 13: Rate of winning trades of a trading game playing the Scattering Spectra (SS) model vs. S&P or the Path-Dependent Volatility (PDV) model vs. S&P. Each heatmap corresponds to a 3 years period, from top to bottom (2015-2017, 2018-2020, 2021-2023).

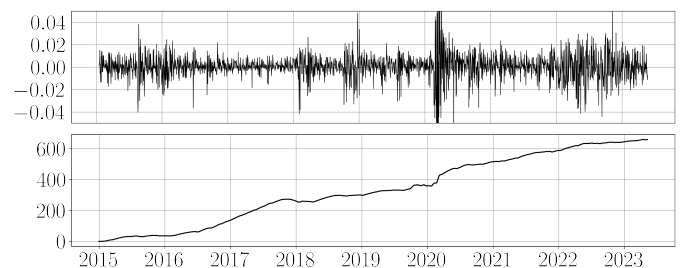


Fig. 14: (Bottom) Aggregated P&L of a trading game playing the Scattering Spectra model vs. the Path-Dependent Volatility model (re-optimized for each  $T$ ), from 2015 to 2023, both calibrated using data from 2000 to 2014. (Top) Time series of S&P log-return on the trading game period.




Fig. 15: P&Ls in a trading game playing the Scattering Spectra model against the Path-Dependent Volatility model. Each heatmap corresponds to a 3 years period, from top to bottom (2015-2017, 2018-2020, 2021-2023).




Fig. 16: Standard deviation of P&Ls in a trading game playing the Scattering Spectra model against the Path-Dependent Volatility (PDV) model. Each heatmap corresponds to a 3 years period, from top to bottom (2015-2017, 2018-2020, 2021-2023).

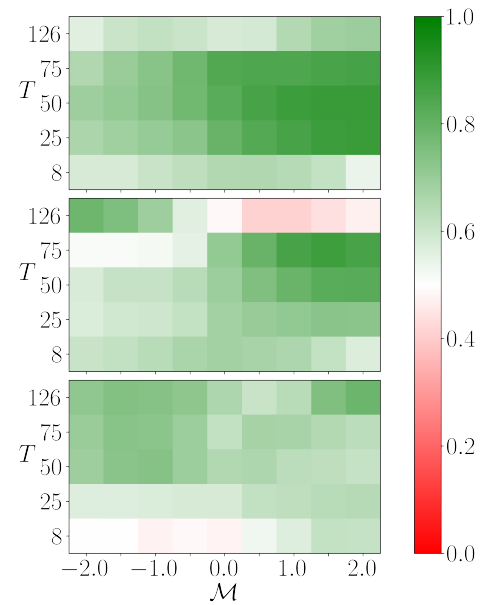


Fig. 17: Percentage of winning trades of a trading game playing the Scattering Spectra model against Path-Dependent Volatility model. Each heatmap corresponds to a 3 years period, from top to bottom (2015-2017, 2018-2020, 2021-2023).