

MINERVA INVESTMENT MANAGEMENT SOCIETY

Quantitative Research Division



Path Dependent Volatility Model

Theoretical and Empirical Study

Ianina Tarca	Co-Head of Quantitative Research
Afonso Tavares	Quantitative Research Analyst
Lorenzo Giambra	Quantitative Research Analyst

Academic Year 2025/2026

Contents

1	Introduction	3
2	From GARCH to Stochastic Volatility	4
3	Path-Dependent Volatility (PDV) Models	7
3.1	Discrete-Time PDV Model	7
3.2	Continuous-Time Formulation of PDV Models	8
3.3	Markovian Approximation of PDV Models	8
4	Empirical Estimation of the PDV Model	10
4.1	Experimental Setup and Implementation	10
4.2	SPX and VIX Time Series	11
4.3	Model Estimation Results	11
4.4	Out-of-Sample Volatility Prediction	12
4.5	Residual Diagnostics	13
4.6	Volatility as a Function of the Trend Feature	13
4.7	The PDV Volatility Surface	14
5	Conclusion	15
A	Simulation and Option Pricing Under the 4-Factor Markov PDV Model	16
A.1	Model Specification	16
A.2	Monte Carlo Simulation Framework	16
A.3	How the PDV Model Enters the Simulation	17
B	Simulation Results and Implied Volatility Smiles	17
B.1	Simulation Setup	18
B.2	SPX Option Prices	18
B.3	SPX Implied Volatility Smile	18
B.4	VIX Implied Volatility Smile	19
B.5	Discussion	20

Abstract

We investigate whether implied volatility can be explained endogenously as a deterministic functional of past asset returns, rather than as the realization of an independent stochastic volatility process. To this end, we study a Path-Dependent Volatility (PDV) model in which volatility is constructed from two observable components: a trend feature capturing leverage effects and an activity feature capturing volatility clustering.

We provide a rigorous theoretical formulation of the PDV framework in both discrete and continuous time and show how long-memory effects can be represented through a Markovian approximation. The model is then estimated empirically using daily data on the S&P 500 index and the VIX over the period 2000–2022.

Our results show that the PDV model explains nearly 87% of the out-of-sample variation in implied volatility using only past index returns. The model successfully reproduces major volatility spikes, including during the COVID-19 crisis. Residual diagnostics confirm statistical robustness, while a geometric analysis reveals that implied volatility lies on a low-dimensional surface spanned by the PDV factors.

These findings provide strong evidence that volatility dynamics are predominantly endogenous and structurally driven by past market behavior. The PDV framework offers a parsimonious and economically interpretable alternative to traditional stochastic volatility models and provides new perspectives for volatility forecasting, risk management, and option pricing.

Disclaimer. This is an academic paper related to an academic project that does not pretend to represent any investment recommendation nor offer any solicitation to buy or sell securities or to adopt an investment strategy. The opinions expressed are subject to change. References to specific securities, asset classes and financial markets are for illustrative purposes only and are not intended to be and should not be interpreted as recommendations. Reliance upon information in this material is at the sole risk and discretion of the reader. The material was prepared only in regard to the specific objectives of Minerva Investment Management Society virtual funds.

© Minerva Investment Management, 2025. All rights reserved.

1 Introduction

Volatility is a central object in modern finance. It governs risk management, portfolio construction, option pricing, and the behaviour of implied volatility surfaces. Market participants treat the VIX as the benchmark measure of future volatility, yet its dynamics remain difficult to reconcile with traditional asset–price models. In particular, Stochastic Volatility (SV) models assume that volatility is driven by its own source of randomness, only partially connected to the price path through the spot–vol correlation parameter, see for instance [11, 2].

This separation between prices and volatility creates persistent modelling challenges. SV models struggle to jointly explain the behaviour of the S&P 500 (SPX) and the VIX, are slow to react to large market moves, and require several latent processes to match empirical regularities. These issues motivate a simpler question: Can we learn the VIX almost entirely from the past path of the underlying index?

The central research question of this paper is therefore: Can volatility be explained and predicted directly from past returns of the SPX, without invoking an independent volatility process? We investigate this through a Path-Dependent Volatility (PDV) model that builds volatility from two intuitive components: the impact of past returns and the intensity of past market movements.

2 From GARCH to Stochastic Volatility

Time series of financial returns exhibit heteroscedasticity (time-varying volatility). Large returns are frequently followed by other large returns, while small returns tend to be followed by small ones, a phenomenon referred to as volatility clustering.

The work of [5] introduced the Autoregressive Conditional Heteroskedasticity (ARCH) model, in which the conditional variance of returns depends on past squared innovations. This framework was subsequently generalized by [4] into the GARCH(p, q) model. The most widely used specification in practice is the GARCH(1, 1) model, defined as:

$$x_t \sim \mathcal{N}(0, \sigma_t^2), \quad \sigma_t^2 = \alpha_0 + \alpha_1 x_{t-1}^2 + \beta_1 \sigma_{t-1}^2. \quad (1)$$

In this formulation, the conditional variance depends on both the most recent squared return innovation and the most recent volatility level. The parameter α_1 captures the short-run impact of return shocks on volatility, while β_1 controls the persistence of volatility over time. Empirically, the sum $\alpha_1 + \beta_1$ is often found to be close to one, implying a high degree of volatility persistence.

Despite their success in capturing volatility clustering, GARCH-type models face important limitations. The conditional variance depends linearly on past squared returns and past variances, which may not adequately reflect the nonlinear feedback mechanisms observed in financial markets. Moreover, positive and negative returns of the same magnitude exert identical effects on volatility, contradicting the leverage effect observed in equity markets. Extensions such as EGARCH and GJR-GARCH introduce asymmetry, but they remain fundamentally parametric. The reliance on fixed functional forms can lead to rigidity, overparameterization, and overfitting.[14, 15]

An alternative is provided by stochastic volatility (SV) models, which replace deterministic variance recursions with a latent stochastic process, modeling return innovations and volatility innovations separately.

Stochastic Volatility models were developed in order to correct the oversimplifying assumption made by the Black-Scholes model that the underlying asset's volatility is constant. Instead, they allow it to fluctuate under a stochastic process, driven by its own sources of randomness, which increases the predictive power of the model relative to the original Black-Scholes equation. A typical Stochastic Volatility (SV) model reads as follows:

$$\frac{dS_t}{S_t} = \sigma_t dW_t, \quad \sigma_t = f(t, Y_t), \quad (2)$$

$$dY_t = \mu(t, Y_t) dt + \nu(t, Y_t) \left(\rho dW_t + \sqrt{1 - \rho^2} dW_t^\perp \right) \quad (3)$$

$$= \mu(t, Y_t) dt + \nu(t, Y_t) dZ_t \quad (4)$$

where $W_t^\perp$ is a Brownian motion independent of W_t ; $\rho \in [-1, 1]$ is the 'spot-vol' correlation (between the price ('spot') and the volatility of the underlying asset); $\sigma_t = f(t, Y_t) > 0$ is the instantaneous volatility; μ is the volatility drift; ν is the volatility of the volatility (vol-of-vol) and Y_t is a positive function that determines the dynamics of implied volatility.

We can thus see that Stochastic Volatility models are driven by two independent Brownian Motions, one of which drives the return process and relates to the implied volatility through the spot-vol correlation ρ , while the other ensures dZ_t is itself a Brownian Motion

and its variance is dt . This new Brownian Motion dZ_t is the source of randomness for the volatility dynamics, and is related to the price's Brownian Motion. Thus, the individual shocks applied to instantaneous volatility reflect those applied to asset returns, although they may adapt slowly to large return shocks.

After integration and using the Brownian Motion of the price process, Y_t is given by:

$$Y_t = Y_0 + \int_0^t \mu(u, Y_u) du + \int_0^t \nu(u, Y_u) \left(\rho \frac{1}{f(u, Y_u)} \frac{dS_u}{S_u} + \sqrt{1 - \rho^2} dW_u^\perp \right). \quad (5)$$

Thus, the implied volatility σ_t at each point may depend on the entire return process of the underlying security $(dS_u/S_u)_{0 \leq u \leq t}$, and three cases can be distinguished depending on the value of ρ :

- $\rho = 0$: the SV model is strictly path-independent with respect to the asset price. The price process and the volatility process are driven by independent Brownian motions ($dZ_t = dW_t^\perp$), and the instantaneous volatility depends only on the volatility driver:

$$\sigma_t = \varphi(t, (W_u^\perp)_{0 \leq u \leq t}).$$

- $\rho \notin \{-1, 0, 1\}$: the SV model is partially path-dependent. The volatility process and the return process are partially correlated, meaning that shocks to returns have a partial feedback effect on volatility:

$$dZ_t = a dW_t + b dW_t^\perp, \quad a, b \neq 0, \quad a^2 + b^2 = 1.$$

In this case, the instantaneous volatility is a function of both the asset and volatility drivers:

$$\sigma_t = \varphi(t, (W_u)_{0 \leq u \leq t}, (W_u^\perp)_{0 \leq u \leq t}).$$

- $\rho \in \{1, -1\}$: the SV model is fully path-dependent. The volatility process and the return process are perfectly correlated (positively or negatively), so that $dZ_t = \pm dW_t$. Volatility can then be written purely as a functional of the asset return (or price) path:

$$\sigma_t = \varphi\left(t, \left(\frac{dS_u}{S_u}\right)_{0 \leq u \leq t}\right) = \psi(t, (S_u)_{0 \leq u \leq t}).$$

However, Stochastic Volatility models have some problems:

- no explanation for mean reversion of volatility
- complicated spot-vol correlation and path dependency when partially path dependent
- functions φ and ψ aren't explicitly chosen, so may take a long time to adapt to large shocks to returns and capture the high volatility associated with them.

Limitations of Classical Stochastic Volatility Models

Despite their flexibility, standard SV models exhibit several structural shortcomings:

- They do not naturally generate strong mean reversion in volatility. Additional parametric constraints are required to force the volatility process back toward its long-run level.
- The spot–volatility correlation introduces an implicit and often unstable form of path dependence. Because this dependence is mediated through correlated Brownian drivers rather than through observable returns, the resulting feedback is difficult to control or interpret.
- The functional forms f , μ , and ν must be specified parametrically. This restricts how quickly volatility can adjust to large, sudden market moves and limits the model’s ability to reproduce extreme or asymmetric volatility dynamics.

These limitations motivate volatility models in which the feedback from prices to volatility is encoded directly through observable returns. Such models avoid latent volatility shocks and provide a clearer, more adaptive representation of how market movements shape volatility dynamics.

3 Path-Dependent Volatility (PDV) Models

Path-Dependent Volatility (PDV) models adopt the opposite perspective to classical SV models. Instead of postulating a separate stochastic driver for volatility, they assume that the volatility at time t is a deterministic functional of the past price (or return) path [9].

In other words, volatility is modeled as a fully endogenous feedback response to past market moves:

$$\sigma_t = \Phi((S_u)_{u \leq t}) \quad \text{or equivalently} \quad \sigma_t = \Psi((r_u)_{u \leq t}),$$

for some measurable functionals Φ or Ψ .

3.1 Discrete-Time PDV Model

In discrete time, letting r_{t-k} denote the (log-)return over day $t-k$, a remarkably effective and empirically tractable specification is a simple linear PDV model of the form

$$\text{Vol}_t = \beta_0 + \beta_1 R_{1,t} + \beta_2 \Sigma_t, \quad \beta_0 > 0, \beta_1 < 0, \beta_2 \in (0, 1), \quad (6)$$

where $R_{1,t}$ is a *trend* feature encoding the impact of past returns, and Σ_t is an *activity* feature encoding past volatility.

The trend feature is defined as a weighted sum of past daily returns,

$$R_{1,t} = \sum_{k \geq 1} w_k^{(1)} r_{t-k}, \quad (7)$$

where the weights $(w_k^{(1)})_{k \geq 1}$ are generated by a time-shifted power-law (TSPL) kernel. Large negative returns make $R_{1,t}$ more negative, thereby increasing Vol_t through (6) when $\beta_1 < 0$. This reproduces the leverage effect: volatility tends to rise after price drops.

The volatility (or activity) feature is built from past squared returns,

$$R_{2,t} = \sum_{k \geq 1} w_k^{(2)} r_{t-k}^2, \quad \Sigma_t = \sqrt{R_{2,t}}, \quad (8)$$

where $(w_k^{(2)})_{k \geq 1}$ are again generated by a TSPL kernel, possibly with different parameters from $(w_k^{(1)})_{k \geq 1}$. This feature captures volatility clustering: large past absolute returns (large r_{t-k}^2) increase Σ_t , and thus, through $\beta_2 > 0$, increase current volatility. The condition $\beta_2 < 1$ plays a role similar to that in GARCH-type models and contributes to the stability of the dynamics.

In practice, the TSPL kernels are chosen so that the weights $w_k^{(1)}$ and $w_k^{(2)}$ are non-negative, sum to one, and decay slowly with the lag k . This slow decay gives the model a long memory of the return path across multiple time scales: very recent returns receive the largest weight, but returns from months or even years ago still have a non-negligible impact, in line with empirical evidence on long memory and rough volatility [10, 6, 1, 7].

Unlike SV models, where feedback from prices to volatility is implicit and mediated through a correlated volatility driver (via the spot-vol correlation ρ), the PDV framework specifies this feedback *explicitly* through the functions $R_{1,t}$ and Σ_t . Volatility is not driven by an additional independent Brownian motion, but is instead treated as an endogenous functional of past returns. This explicit path dependence allows the model to react quickly and consistently to large return shocks, to capture leverage and volatility clustering in a transparent way, and to remain simple, interpretable, and empirically testable.

3.2 Continuous-Time Formulation of PDV Models

While the PDV model (6) is formulated in discrete time for empirical estimation, it naturally admits a continuous-time analogue that defines a full diffusion model for asset prices. Let $(W_t)_{t \geq 0}$ be a standard Brownian motion and consider the price dynamics

$$\frac{dS_t}{S_t} = \sigma_t dW_t, \quad (9)$$

where the instantaneous volatility is specified as the path-dependent functional

$$\sigma_t = \beta_0 + \beta_1 R_{1,t} + \beta_2 \Sigma_t. \quad (10)$$

The continuous-time counterparts of the discrete features (7)–(8) are defined through convolution integrals:

$$R_{1,t} = \int_0^t K_1(t-u) d \log S_u, \quad (11)$$

$$R_{2,t} = \int_0^t K_2(t-u) d \langle \log S \rangle_u, \quad \Sigma_t = \sqrt{R_{2,t}}, \quad (12)$$

where K_1 and K_2 are time-shifted power-law (TSPL) kernels that encode both short- and long-memory effects, and $\langle \log S \rangle_t$ denotes the quadratic variation of the log-price process.

Under this formulation, volatility is no longer an independent stochastic process driven by its own Brownian shock, as in classical SV models. Instead, it evolves as a nonlinear functional of the entire past return path. The resulting PDV diffusion is intrinsically non-Markovian: the future evolution of (S_t, σ_t) depends on the full history of past log-returns through $(R_{1,t}, \Sigma_t)$.

This structure highlights the conceptual shift relative to SV models. In SV settings, volatility follows an exogenous stochastic differential equation, possibly correlated with the price. In PDV models, by contrast,

- volatility has no independent source of randomness,
- all randomness enters through the price Brownian motion,
- the feedback from past prices to volatility is explicit and deterministic.

As a result, PDV models are pure feedback systems in which volatility is fully endogenous.

3.3 Markovian Approximation of PDV Models

PDV models are inherently non-Markovian because the TSPL kernels K_1 and K_2 introduce long-memory effects. A tractable approximation is obtained by replacing these kernels with finite mixtures of exponential functions. Specifically, one approximates

$$K_j(\tau) \approx \sum_{m=1}^{M_j} a_{j,m} \lambda_{j,m} e^{-\lambda_{j,m} \tau}, \quad j = 1, 2, \quad (13)$$

where $a_{j,m} \geq 0$, $\sum_m a_{j,m} = 1$, and $\lambda_{j,m} > 0$ determine the short- and long-term decay structure.

Under this representation, the path-dependent features $(R_{1,t}, R_{2,t})$ become weighted sums of auxiliary Ornstein–Uhlenbeck state variables satisfying

$$dX_{m,t}^{(1)} = -\lambda_{1,m}X_{m,t}^{(1)}dt + d\log S_t, \quad (14)$$

$$dX_{m,t}^{(2)} = -\lambda_{2,m}X_{m,t}^{(2)}dt + d\langle \log S \rangle_t. \quad (15)$$

The PDV features are then recovered through

$$R_{1,t} = \sum_{m=1}^{M_1} a_{1,m}X_{m,t}^{(1)}, \quad R_{2,t} = \sum_{m=1}^{M_2} a_{2,m}X_{m,t}^{(2)}.$$

The joint process

$$\left(S_t, X_{1,t}^{(1)}, \dots, X_{M_1,t}^{(1)}, X_{1,t}^{(2)}, \dots, X_{M_2,t}^{(2)} \right)$$

is therefore Markovian, yet reproduces the long-memory behaviour of the original PDV model. This Markovian approximation enables efficient simulation, facilitates option pricing, and allows direct comparison with standard stochastic volatility models within a unified diffusion framework.

4 Empirical Estimation of the PDV Model

In this section we take the PDV model defined in Section 3 as given and focus exclusively on its empirical estimation from market data. A publicly available implementation of related PDV computations is documented in [12]. The goal is to assess to what extent the volatility of the S&P 500 can be explained directly from the past path of index returns.

4.1 Experimental Setup and Implementation

All empirical experiments are implemented in Python. Data are retrieved from Yahoo! Finance using the `yfinance` package. We consider daily observations of the S&P 500 index (ticker `^GSPC`) and the VIX index (ticker `^VIX`).

The sample period starts on January 1, 1995, which provides a sufficiently long burn-in window for computing historical convolution features. The training period spans from January 1, 2000 to December 31, 2018, while the out-of-sample test period runs from January 1, 2019 to May 15, 2022. This split allows us to evaluate the predictive performance of the PDV model during both tranquil market conditions and major stress episodes such as the COVID-19 crisis.

Daily closing prices are converted into log-returns according to

$$r_t = \log \left(\frac{S_t}{S_{t-1}} \right).$$

The VIX index is scaled to decimal form before estimation.

For the computation of the PDV features $R_{n,t}$, the maximum memory length is fixed to $\Delta_{\max} = 1000$ business days. The kernels are chosen either as time-shifted power-law (TSPL) kernels or as convex combinations of two exponential kernels. In our baseline specification, we set

$$p = 1, \quad \text{setting} = \{(1, 1), (2, 1/2)\},$$

which corresponds to the linear PDV model

$$\text{Vol}_t = \beta_0 + \beta_1 R_{1,t} + \beta_2 \Sigma_t.$$

All estimation routines are wrapped into a unified function `perform_empirical_study`, which computes the PDV features, fits the model on the training sample, and evaluates it on the test sample.

4.2 SPX and VIX Time Series




Figure 1: Joint evolution of the S&P 500 (SPX) and the VIX index over the sample period.

Figure 1 displays the joint evolution of the S&P 500 index and the VIX over the full sample period. The SPX exhibits a long-term upward trend punctuated by sharp drawdowns, most notably during the 2008 Global Financial Crisis and the 2020 COVID-19 crisis. In contrast, the VIX is mean-reverting but displays strong, abrupt spikes during periods of market stress.

This joint dynamics highlights the pronounced leverage effect: volatility rises sharply following negative market returns. These features motivate the PDV modeling approach, which aims to explain volatility endogenously from the past return path.

Parameter	Estimated Value
β_0	0.0507
β_1	-0.0798
β_2	0.8617
θ_1	0.7994
θ_2	0.4239
$\lambda_1^{(1)}, \lambda_1^{(2)}$	50.77, 16.99
$\lambda_2^{(1)}, \lambda_2^{(2)}$	3.80, 1.20

Table 1: Estimated PDV parameters on the training sample. The coefficients $\beta_1 < 0$ and $\beta_2 > 0$ capture leverage and volatility clustering, respectively. The parameters θ_j and $\lambda_j^{(m)}$ define the two-scale exponential kernel approximations used in the Markovian representation of the PDV features.

4.3 Model Estimation Results

The PDV model is estimated on the training sample by least-squares regression. The resulting optimal parameters are reported in Table 1. The signs of the estimated coef-

ficients are consistent with financial theory: the trend coefficient is negative, capturing leverage effects, while the activity coefficient is positive, capturing volatility clustering.

These results confirm that a large fraction of implied volatility dynamics can be explained directly from the past return path without introducing an independent volatility driver.

Table 2 reports the in-sample and out-of-sample performance of the PDV model. The model achieves a very high in-sample R^2 of 94.8% and retains strong predictive power out-of-sample with a test R^2 of 86.9%. This confirms that a large fraction of implied volatility dynamics can be explained using only past SPX returns.

Metric	Training	Test
R^2	0.948	0.869
RMSE	0.0196	0.0334

Table 2: In-sample and out-of-sample performance of the PDV model.

4.4 Out-of-Sample Volatility Prediction

Figure 2 compares the predicted volatility produced by the PDV model with the observed VIX on the out-of-sample test period from 2019 to 2022. The model successfully tracks both the level and the major spikes in volatility, including the extreme surge observed during the COVID-19 market crash of March 2020.

The PDV model reacts sharply to large negative returns and elevated market activity, demonstrating its ability to generate fast endogenous volatility responses without relying on an independent stochastic volatility driver. While some short-term overshooting and smoothing effects are visible, the overall dynamics of implied volatility are remarkably well captured.




Figure 2: Out-of-sample prediction of the VIX using the PDV model (red) versus observed VIX (black). The dashed blue curve represents the S&P 500 index.

The empirical results strongly support the hypothesis that volatility is predominantly path-dependent. The PDV model explains nearly 87% of the out-of-sample variance of the VIX using only past SPX returns. This suggests that most volatility dynamics

are endogenous and structurally driven by past market movements rather than by an independent volatility shock.

These findings provide strong empirical backing for the PDV modeling paradigm and help explain why traditional stochastic volatility models often require extreme spot-vol correlations and multiple latent factors to match observed implied volatility dynamics.

4.5 Residual Diagnostics

Figure 3 presents a scatter plot of predicted versus observed volatility on the test set together with the empirical distribution of residuals. The points concentrate closely around the diagonal reference line, indicating a strong agreement between predicted and true values.

The residual distribution is approximately symmetric and centered around zero, suggesting the absence of systematic bias in the PDV predictions. A moderate degree of heteroskedasticity remains at higher volatility levels, which is expected given the extreme behavior of volatility during crisis periods. Overall, the residual analysis confirms that the PDV model provides an unbiased and statistically well-behaved approximation of implied volatility.




Figure 3: Predicted versus observed volatility on the test set with residual distribution.

4.6 Volatility as a Function of the Trend Feature

Figure 4 shows the empirical relationship between volatility and the trend feature $R_{1,t}$ on the training set. A clear negative dependence emerges: large negative values of $R_{1,t}$, corresponding to sustained negative returns, are associated with high levels of implied volatility.

This monotonic decreasing structure provides a direct, nonparametric visualization of the leverage effect. The fitted PDV regression curve superimposed on the scatter plot captures this relationship remarkably well, confirming that the linear dependence in $R_{1,t}$ is sufficient to encode the dominant leverage mechanism.




Figure 4: Implied volatility as a function of the trend feature R_1 . Color scale represents Σ_t .

4.7 The PDV Volatility Surface

Figure 5 displays the three-dimensional relationship between the PDV features ($R_{1,t}, \Sigma_t$) and the implied volatility level. The figure reveals that implied volatility lies on a well-structured surface that is approximately linear in the two PDV factors.

The trend feature $R_{1,t}$ controls the asymmetry through the leverage channel, while the activity feature Σ_t controls the overall volatility level through volatility clustering. The nearly planar geometry of the surface confirms that the two-factor linear PDV model provides an accurate low-dimensional representation of implied volatility dynamics.




Figure 5: Three-dimensional PDV volatility surface: implied volatility as a function of (R_1, Σ) .

5 Conclusion

This paper set out to test a fundamental modeling question in volatility dynamics: whether implied volatility is primarily driven by an independent stochastic volatility process, as assumed in classical stochastic volatility (SV) models, or whether it can instead be explained endogenously as a deterministic functional of the past return path of the underlying asset.

We addressed this question by developing and estimating a Path-Dependent Volatility (PDV) model in which volatility is constructed explicitly from two observable components: a trend feature capturing leverage effects and an activity feature capturing volatility clustering. The PDV model was estimated on the S&P 500 and evaluated against the VIX over a long out-of-sample test period including the COVID-19 market crash.

Our empirical results provide strong support for the PDV paradigm. The model explains nearly 87% of the out-of-sample variation in implied volatility using only past index returns. This level of explanatory power is remarkable given that no independent volatility shock is introduced. The residual diagnostics confirm that the model is statistically well behaved, while the geometric analysis reveals that implied volatility lies on a low-dimensional surface spanned by the PDV factors.

These findings suggest that volatility is predominantly endogenous and structurally driven by past market dynamics rather than by an independent source of randomness. From a modeling perspective, this offers a conceptual explanation for why classical stochastic volatility models often require extreme spot–vol correlations and multiple latent factors to match observed implied volatility dynamics.

The PDV framework offers a parsimonious and economically interpretable alternative to traditional volatility modeling. By treating volatility as an emergent feedback effect of past price dynamics, it provides a unified structural interpretation of leverage, volatility clustering, and regime shifts. This perspective opens new directions for volatility forecasting, risk management, and option pricing that do not rely on latent volatility shocks, but instead treat volatility as an endogenous property of the market.

A Simulation and Option Pricing Under the 4-Factor Markov PDV Model

This appendix illustrates how the PDV framework can be embedded into a low-dimensional Markovian diffusion suitable for Monte Carlo simulation and option pricing. The construction follows the 4-factor Markov PDV model of [8], which provides a tractable Markovian approximation to the non-Markovian PDV dynamics introduced in Section 3.[13]

A.1 Model Specification

Under the risk-neutral measure, the underlying asset price satisfies

$$\frac{dS_t}{S_t} = \sigma_t dW_t, \quad (16)$$

where the instantaneous volatility is given by the linear PDV specification

$$\sigma_t = \beta_0 + \beta_1 R_{1,t} + \beta_2 \sqrt{R_{2,t}}. \quad (17)$$

Each PDV feature is represented by a convex mixture of two exponential kernels:

$$R_{n,t} = (1 - \theta_n) R_{n,0,t} + \theta_n R_{n,1,t}, \quad n \in \{1, 2\}. \quad (18)$$

The four latent Markovian factors evolve according to

$$dR_{1,j,t} = \lambda_{1,j} (\sigma_t dW_t - R_{1,j,t} dt), \quad j \in \{0, 1\}, \quad (19)$$

$$dR_{2,j,t} = \lambda_{2,j} (\sigma_t^2 - R_{2,j,t}) dt, \quad j \in \{0, 1\}, \quad (20)$$

with decay rates $\lambda_{n,j} > 0$ and mixture parameters $\theta_n \in (0, 1)$.

The resulting state vector

$$(S_t, R_{1,0,t}, R_{1,1,t}, R_{2,0,t}, R_{2,1,t})$$

is Markovian, enabling efficient simulation while preserving the long-memory effects of the original PDV convolution kernels.

A.2 Monte Carlo Simulation Framework

The 4-factor Markov PDV model is implemented in Python using `PyTorch`, allowing efficient GPU-accelerated Monte Carlo simulation. The procedure is as follows:

1. **Data Initialisation.** Daily SPX prices are downloaded via `yfinance`. The previous four years of returns are used to initialise the state variables $R_{n,j,t}$ through empirical convolution approximations.
2. **Parameter Specification.** The model parameters $(\beta_0, \beta_1, \beta_2)$ and $(\theta_1, \theta_2, \lambda_{1,0}, \lambda_{1,1}, \lambda_{2,0}, \lambda_{2,1})$ are taken from the empirical estimation in Section 4.
3. **Path Simulation.** The SDE system is discretised using the Euler–Maruyama scheme, implemented through the `TorchMonteCarloExponentialModel` class, which handles vectorised and GPU-parallel simulation of all paths.

4. **Option Pricing.** European call prices are computed under the risk-neutral measure as

$$C(T, K) = \mathbb{E}[(S_T - K)^+],$$

and implied volatilities are obtained by numerical inversion of the Black-Scholes formula.

A.3 How the PDV Model Enters the Simulation

The Markovian PDV model replaces the stochastic volatility process of classical SV models with an endogenous volatility functional driven entirely by the simulated return path. At each Monte Carlo time step, the simulation proceeds as follows:

1. **Update the PDV factors.** The latent factors $R_{1,0,t}$, $R_{1,1,t}$, $R_{2,0,t}$, $R_{2,1,t}$ are updated using the Markovian SDE system

$$dR_{1,j,t} = \lambda_{1,j}(\sigma_t dW_t - R_{1,j,t}dt), \quad dR_{2,j,t} = \lambda_{2,j}(\sigma_t^2 - R_{2,j,t})dt.$$

These factors approximate the PDV convolution features from Section 3.

2. **Reconstruct the PDV features.** The trend and activity features are computed as

$$R_{1,t} = (1 - \theta_1)R_{1,0,t} + \theta_1 R_{1,1,t}, \quad R_{2,t} = (1 - \theta_2)R_{2,0,t} + \theta_2 R_{2,1,t}.$$

3. **Compute instantaneous volatility.** Volatility is given by the PDV formula

$$\sigma_t = \beta_0 + \beta_1 R_{1,t} + \beta_2 \sqrt{R_{2,t}},$$

so that all randomness in σ_t arises endogenously from the simulated return path.

4. **Simulate the next price.** The asset evolves according to

$$S_{t+\Delta t} = S_t \exp\left(-\frac{1}{2}\sigma_t^2 \Delta t + \sigma_t \sqrt{\Delta t} Z_t\right),$$

where $Z_t \sim \mathcal{N}(0, 1)$.

This feedback loop makes the PDV diffusion a self-exciting system in which volatility reacts instantaneously to past returns, generating realistic leverage and volatility clustering effects without the need for an additional stochastic volatility driver. The same mechanism applies to nested simulations used for VIX option pricing.

B Simulation Results and Implied Volatility Smiles

We now present the numerical results obtained from Monte Carlo simulation of the 4-factor Markov PDV model. The purpose is to assess the model's capacity to generate realistic option prices and implied volatility smiles for both SPX and VIX derivatives.

B.1 Simulation Setup

The Markovian PDV process is simulated under the risk-neutral dynamics described above. The PDV factors are initialised from historical SPX returns using a four-year lookback window.

Model parameters used in the simulation are:

- exponential decay rates $\lambda_1 = (55, 101)$, $\lambda_2 = (20, 3)$,
- PDV regression coefficients $\beta_0 = 0.04$, $\beta_1 = -0.13$, $\beta_2 = 0.65$,
- mixture parameters $\theta_1 = 0.25$, $\theta_2 = 0.50$.

We simulate:

$N = 10,000$ price paths, $N_{\text{VIX}} = 1,000$ nested paths for VIX computation,

using 10 time steps per business day and a maturity of $T = 0.25$ years.

B.2 SPX Option Prices

Figure 6 shows the simulated call price curve. The curve is monotonically decreasing and convex, consistent with standard no-arbitrage properties. The forward price is indicated by a vertical dashed line.




Figure 6: Simulated SPX call option prices under the 4-factor Markov PDV model.

B.3 SPX Implied Volatility Smile

The implied volatility smile in Figure 7 exhibits the characteristic features of equity index options:

- a downward-sloping skew,
- elevated implied volatilities for out-of-the-money puts,

- a smooth and nearly linear structure.

This demonstrates that the PDV model can reproduce realistic SPX smiles using only two endogenous PDV factors.




Figure 7: SPX implied volatility smile generated by the PDV model.

B.4 VIX Implied Volatility Smile

A key advantage of the PDV framework is that VIX dynamics arise endogenously from the return path, allowing simultaneous simulation of SPX and VIX derivatives. Figure 8 displays the VIX implied volatility smile produced by nested simulation.[16]

The smile displays the typical upward-sloping shape observed in VIX markets, reflecting the convexity of the VIX payoff and the asymmetric impact of large volatility shocks.[3]




Figure 8: VIX implied volatility smile generated by the PDV model.

B.5 Discussion

These simulation results demonstrate that the 4-factor Markovian PDV model:

- generates realistic SPX option prices and implied volatility smiles,
- reproduces both the SPX skew and the VIX upward-sloping smile,
- achieves this without introducing an independent stochastic volatility driver,
- relies purely on deterministic feedback from past returns.

These findings reinforce the central message of the paper: a large part of the structure of implied volatility surfaces appears to be *endogenous*, arising naturally from the return path of the underlying asset. PDV models therefore offer a parsimonious and economically interpretable alternative to classical stochastic volatility models for pricing equity and volatility derivatives.

References

- [1] Eduardo Abi Jaber and Omar El Euch. Markovian structure of the volterra heston model. *SIAM Journal on Financial Mathematics*, 2019.
- [2] L. Bergomi. *Stochastic Volatility Modeling*. Financial Mathematics Series. Chapman & Hall/CRC, 2016.
- [3] Lorenzo Bergomi and Julien Guyon. The smile in stochastic volatility models. *Risk Magazine*, 2012.
- [4] Tim Bollerslev. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3):307–327, 1986.
- [5] Robert F. Engle. Autoregressive conditional heteroskedasticity with estimates of the variance of united kingdom inflation. *Econometrica*, 50(4):987–1007, 1982.
- [6] J. Gatheral, T. Jaisson, and M. Rosenbaum. Volatility is rough. *Quantitative Finance*, 18(6):933–949, 2018.
- [7] J. Gatheral, P. Jusselin, and M. Rosenbaum. The quadratic rough heston model and the joint calibration problem. *Risk Magazine*, May 2020.
- [8] Julien Guyon. A 4-factor markovian model for path-dependent volatility. *SSRN Working Paper*, 2025.
- [9] Julien Guyon and Jordan Lekeufack. Volatility is (mostly) path-dependent. *SSRN Electronic Journal*, 2023. Bloomberg L.P. Quantitative Research and UC Berkeley, Department of Statistics.
- [10] A. Harvey. Long memory in stochastic volatility. In J. Knight and S. Satchell, editors, *Forecasting Volatility in Financial Markets*, pages 307–320. Butterworth-Heinemann, 1998.
- [11] D. Hobson and L. C. G. Rogers. Complete models with stochastic volatility. *Mathematical Finance*, 8(1):27–48, 1998.
- [12] Jordy Lekeufack. Volatility is mostly path dependent. <https://github.com/Jordylek/VolatilityIsMostlyPathDependent>, 2022.
- [13] Robert C. Merton. Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics*, 1976.
- [14] Daniel Nelson. Conditional heteroskedasticity in asset returns: A new approach. *Econometrica*, 1991.
- [15] Yasuhiro Omori, Siddhartha Chib, Neil Shephard, and Jouchi Nakajima. Stochastic volatility with leverage: Fast likelihood inference. *Journal of Econometrics*, 2007.
- [16] Artur Sepp. Pricing vix options in a jump-diffusion model. *Wilmott Magazine*, 2011.