A Markov Interpretation of Rough Volatility

Gregory Pelts

November 1, 2026
Research preprint

Abstract

An equity-volatility model must explain how the smile moves, as well as where it stands today. We propose a Markov interpretation of rough volatility in which a pair of interdependent financial clocks connects a compact market state to activity at different horizons. The emphasis is on observable consequences: skew stickiness, maturity scaling, persistent volatility and the forward smile. An SPXW prototype fits a broad option cross-section without a separately fitted, state-dependent leverage function. A comparison with Heston calibrated to the same quotes shows that both models generate downside skew and curved call wings, while implying materially different prices for future-reset exposure. The prototype's forward smiles retain the character of the short-tenor spot skew, consistent with a practitioner expectation of persistent smile shape. Published local- and stochastic-local-volatility comparisons place these results in context. The evidence supports a practical pricing framework and identifies tests of its dynamics; vanilla agreement alone does not establish the accuracy of exotic prices or hedges.

Keywords: rough volatility; Markov models; equity options; smile dynamics; skew-stickiness ratio; forward-start options.

AI-use disclosure. OpenAI tools assisted manuscript drafting, editing, reference checking and development of the research software.

1 The smile is a moving risk exposure

A trader buying downside protection is exposed not only to the current price of a put, but also to how that put will reprice after an index move. For a market maker, the relevant questions include whether the smile steepens in a sell-off, how much at-the-money volatility responds, and whether the response persists at longer maturities. For a structured-product desk, the distribution of the smile at a future reset date can matter as much as today's smile.

These questions concern dynamics. Agreement with today's European option prices does not uniquely determine tomorrow's surface or the value of a claim that depends on prices at several dates. The forward-smile bounds studied by Badikov, Jacquier, Liu and Roome [12] make this distinction precise. It is therefore possible for models with similar vanilla fits to imply different prices and hedges for future-reset exposure.

Our starting point is to design the dynamics around established market regularities. The objective is a model whose short-dated skew, response to spot, and longer-horizon behaviour can be assessed together. A compact Markov description is useful here because it gives a current risk state from which to value future scenarios. It need not require volatility to have the same apparent smoothness at every observed horizon.

2 What the market asks the model to explain

2.1 Roughness and maturity scaling

Gatheral, Jaisson and Rosenbaum [1] find empirical evidence of highly irregular log-volatility at short horizons, with estimated roughness well below the Brownian benchmark. Bayer, Friz and Gatheral [2] connect rough volatility to equity option pricing. Fukasawa [5] examines the implications of short-maturity skew scaling for volatility regularity. A central market signature is the rapid increase in the magnitude of the at-the-money skew as maturity shortens. A familiar rough-model scaling is

$$|\partial_k \sigma_{\mathrm{imp}}(0,T)| \;\sim\; T^{\mathcal{H}-1/2}, \qquad k=\log(K/F_T), \qquad 0<\mathcal{H}<\tfrac{1}{2}. \tag{1}$$

Here T is time to expiry, F_T the forward, and $\mathcal{H}$ a roughness exponent. The expression describes a slope of implied volatility, not a rule for the volatility level itself. Its validity for a particular model and market horizon must be checked.

For a desk, the practical question is how the skew changes when moving from daily to weekly and monthly options. Matching one short-dated smile is weaker evidence than reproducing that progression. Short-horizon roughness also should not dictate all long-horizon behaviour: Bennedsen, Lunde and Pakkanen [3] distinguish the irregularity of local fluctuations from the persistence of volatility over longer periods.

2.2 Skew stickiness: the smile after a spot move

The *skew-stickiness ratio* (SSR), developed in Bergomi's work on smile dynamics [6], compares the response of at-the-money implied volatility to an index return with the slope already visible in the smile. In a fixed-maturity empirical measurement, write schematically

$$\Delta\sigma_{\mathrm{ATM}}(T) = \beta(T)\,\Delta\log S + \text{residual}, \qquad \mathrm{SSR}(T) = \frac{\beta(T)}{\partial_k \sigma_{\mathrm{imp}}(0,T)}. \tag{2}$$

Returns and volatility changes must use a common observation interval; roll-down and carry conventions must also be kept consistent. The ratio is informative where the skew is materially different from zero. In an equity sell-off, a negative spot/volatility response and a negative skew normally give a positive ratio.

SSR is directly relevant to the volatility component of a delta hedge. Friz and Gatheral [4] derive a short-maturity limit of two in the regular diffusion setting and $\mathcal{H}+3/2$ in the rough setting they study. These are useful reference points, not identities imposed on every model with irregular volatility. They also show that the forward-variance curve affects SSR, so its behaviour is richer than a single constant correlation.

An empirical comparison remains essential. Bourgey, Delemotte and De Marco [8] find that, after fitting the SPX volatility term structure, rough and classical forward-variance models can have similar SSRs. Roughness alone therefore does not establish better smile dynamics. The relevant test is the joint maturity profile of skew, SSR and smile curvature, and how that profile changes across market states.

2.3 Financial activity and the passage of time

A trading session, an overnight interval and a weekend need not carry the same amount of variance per calendar hour. The broader idea of modelling price evolution in a clock related to financial activity

goes back to Clark [9]. It provides a natural language for separating elapsed time from the amount of risk accumulated during that time.

Two distinctions matter in practice. The first is the variance term structure across contractual maturities. The second is the stochastic pace at which the market changes state. Fitting the former does not prove that the latter is correct. A useful model should accommodate both and then face a dynamic test, rather than rely on a fitted term structure to explain every change in the surface.

3 A compact interpretation of rough volatility

The earlier *Roughly Rough* paper [10] develops the financial-clock viewpoint. The present equity programme uses a pair of interdependent clocks: the market state influences the pace of activity, and that activity feeds back into volatility. Different horizons can then experience different effective rates of evolution.

The organizing idea can be summarized without specifying the internal construction:

$$\text{current market state} \longleftrightarrow \text{interdependent financial clocks} \longrightarrow \text{joint spot and volatility dynamics.}$$

A Markov interpretation means that the current state carries the information used to project future scenarios. It does not mean that all observable volatility changes must resemble a smooth, constant-speed diffusion. Nor is a fractional noise specification the only possible route to rough-looking market behaviour.

The intended advantage is flexibility in the relationship between short and long horizons while retaining a compact state for pricing and risk. The equity extension is designed to let the smile level, skew and curvature respond to that state. The economic requirement is stronger than a good cross-sectional fit: the same structural description should continue to explain the surface after the market has moved.

The smile is generated by these dynamics rather than by a separately calibrated function multiplying volatility at every time and spot level. This preserves a direct route from the specified risk states to future option values. It is a modelling choice whose consequences can be tested, not a guarantee that every observed smile dynamic will be reproduced.

Table 1: Market observations and the tests they imply. These are validation criteria, not claims that the prototype has already passed every test.

Market feature	Relevant test
Steep short-dated skew	Reproduce the change in skew across daily, weekly and monthly expirations.
Spot/volatility leverage	Match SSR and smile-shape changes after positive and negative index returns.
Roughness with persistence	Separate short-horizon scaling from the decay of longer-horizon volatility dependence.
Uneven accumulation of risk	Explain maturity and trading-calendar effects without confusing fitted variance levels with validated dynamics.
Exposure to future smiles	Compare forward-start prices and hedge behaviour under a common vanilla-market fit.

4 Evidence from the equity prototype

The SPXW study uses ThetaData end-of-day data for October 30, 2026 and 21 expirations from November 2 to December 1, spanning approximately three to 32 calendar days. It fits 5,551 positive-volume

quotes, including both calls and puts and both in- and out-of-the-money options. The targets are the reported closing prices; the observed bid/ask spreads and trading volumes determine their calibration weights, following Pelts [11].

The study uses one common set of dynamic coefficients with a separately fitted maturity profile. Figure 1 illustrates the agreement between model and market smile levels and shapes at six maturities spanning the window, with Heston fitted to the same quotes shown alongside. These are ordinary options with strikes fixed today. All 21 expirations are included in the calibration, including those not illustrated.

This is evidence of practical vanilla-price fit. The distinction between that result and a test of market dynamics is deliberate: the snapshot does not yet establish SSR agreement, historical state-transition accuracy or an advantage in forward-start pricing. Those are the next comparisons that can distinguish the interpretation from another way of fitting a surface.

5 A practical test: the value of a future smile

5.1 Forward-start options as a benchmark

A forward-start option fixes its strike at a future date. For example, a cash-settled forward-start straddle pays

$$|S_{T_2} - K S_{T_1}|, \qquad 0 < T_1 < T_2, \tag{3}$$

where K is the relative strike. At $K = 1$, it measures the magnitude of the move after the reset. Unlike an ordinary European option, it depends on the joint distribution at two dates. It is therefore a clean example of exposure that a fit to separate vanilla maturities does not determine.

Such exposures also arise within cliquets and other reset products, which motivate Bergomi's *Smile Dynamics* programme [7]. The economic question is the value of volatility encountered after the market has changed state.

An important practitioner expectation is that a forward smile should retain the recognizable character of today's smile at a comparable remaining tenor. Albanese, Lo and Mijatović [13, Section 5.2] explicitly use this criterion when liquid forward-start quotes are unavailable: shifting the reset date should not, by itself, erase the characteristic smile shape. For an equity desk, persistence of downside skew is therefore a meaningful qualitative test, alongside the fit to today's vanilla options.

5.2 A numerical example from the literature

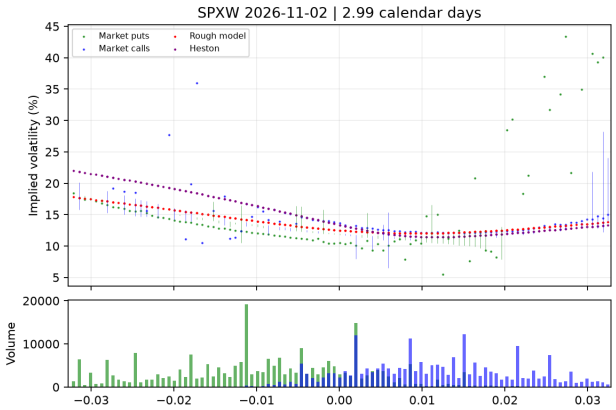
Badikov et al. [12, Section 5.1, Figure 2] consider a one-year reset followed by a six-month option period. Both terminal distributions agree with Black–Scholes volatility of 20%, with spot normalized to one and zero carry. Yet the same two vanilla marginals admit a much lower at-the-money forward-start value, as summarized in Table 2.

Table 2: Published annualized at-the-money forward-volatility example: one-year reset, eighteen-month maturity. Values from Badikov et al. [12], Section 5.1, pp. 11–12.

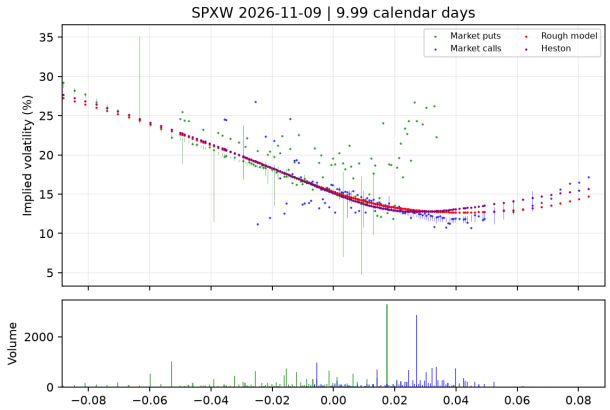
Dynamics or bound consistent with the vanilla marginals	Forward implied volatility
Black–Scholes dynamics	20.00%
Hobson–Klimmek lower bound	6.95%
Numerical linear-programming lower bound	6.98%

The lower-bound rows are two calculations of the same extremal problem, not observed market

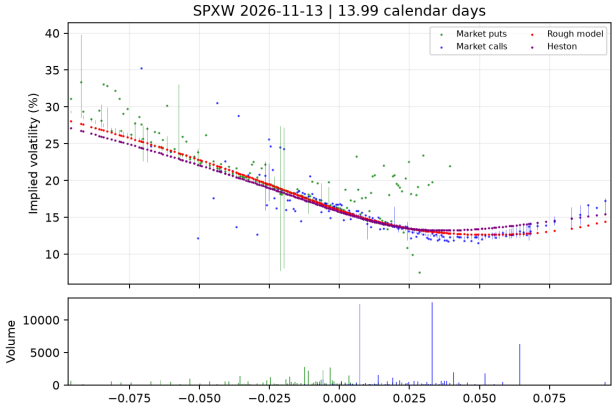
November 2: approximately three days



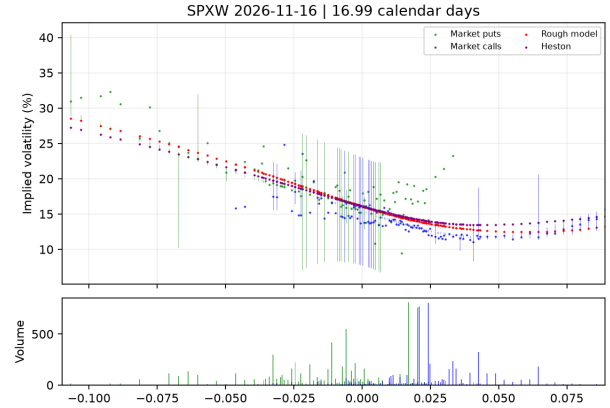
November 9: approximately ten days



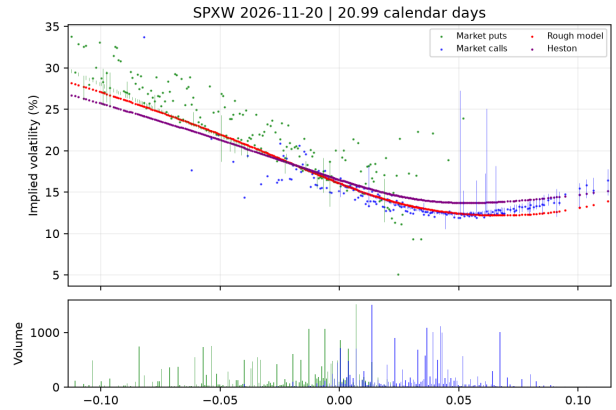
November 13: approximately two weeks



November 16: approximately seventeen days



November 20: approximately three weeks



December 1: approximately one month

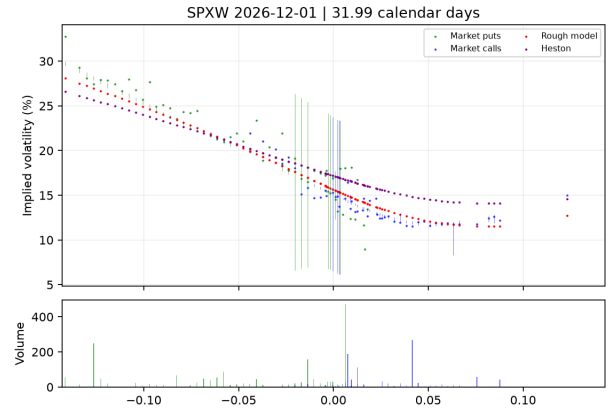


Figure 1: SPXW implied-volatility smiles from the October 30, 2026 snapshot. Green and blue points are observed vanilla put and call closing prices expressed as implied volatilities; red points are the clock-based model and purple points are Heston. The horizontal axis is $\log(K/F)$. Vertical ranges show observed bid/ask implied volatilities; lower panels show volume. These are vanilla market data, not market forward-start smiles. The clock-based values use 524,288 Monte Carlo paths. Data source: ThetaData.

quotes. This example quantifies how much the joint dynamics can matter even when vanilla prices at the two dates agree.

5.3 Local versus stochastic volatility: a matched-market comparison

Van der Stoep, Grzelak and Oosterlee [15] compare local volatility, Heston and Heston stochastic-local volatility against common vanilla-market data. Their Table 2 documents the vanilla fits; Figures 6 and 7, on p. 1450045-21, compare forward smiles for two-year option periods beginning in years two and six.

In the well-calibrated Heston case, its forward smile and that of the stochastic-local model are close. The local-volatility smile is much flatter, with a substantial difference around the money despite the good vanilla fit. The difference persists for the later reset. When the initial Heston fit is poorer, the stochastic-local correction restores the vanilla fit and produces a forward smile between the pure local and stochastic descriptions. These are model comparisons, not a comparison with traded forward-start quotes; they identify a concrete risk that vanilla fit alone does not resolve.

5.4 The leverage adjustment and the meaning of volatility parameters

Stochastic-local volatility restores the vanilla fit through an additional function of time and spot [15]. This leverage function changes how the stochastic-volatility state translates into actual spot variance. The original volatility parameters therefore lose their direct interpretation as parameters of that variance, although they retain their meaning within the latent variance process. Forward-start prices depend on both the variance dynamics and the fitted leverage function, including its behaviour outside well-constrained market states. This introduces a further source of model risk: two different stochastic-volatility specifications can each be corrected to today's surface without agreeing on future smiles. Sensitivity to this choice matters for a desk; instability in a particular implementation must be established by tests.

The clock-based prototype fits the smile through its specified state dynamics, without a separately fitted, state-dependent leverage function. Its maturity profile adjusts activity across horizons, rather than adding a separate spot-dependent volatility multiplier. The intended advantage is a more direct connection between the model's risk states and its predictions for future volatility. The Heston comparison below examines those predictions; market observations are still needed to establish their accuracy.

5.5 Forward-start prices from the equity prototype

The SPXW calibration in Figure 1 also generates prices for options whose strikes are fixed later. We retain its coefficients and maturity profile and consider three reset dates, each followed by an option period of approximately two weeks. No market forward-start quotes are available in this dataset. The curves below are model predictions; the observed spot-start market smile is shown directly alongside them in Figure 2. The spot reference is the November 13 expiry, observed on October 30, with 13.99 days remaining. This is the nearest quoted tenor to the three forward option periods of 14, 14 and 15 days. The last comparison therefore has a one-day tenor mismatch.

To separate smile shape from deterministic carry, moneyness is measured using today's forward curve: a relative strike K corresponds to the cash payoff $(S_{T_2} - KS_{T_1}F_0(T_2)/F_0(T_1))^+$, where F_0 is today's forward curve. Implied volatility is annualized over the period from reset to maturity. The benchmark is the standard Heston stochastic-volatility model [14], fitted independently to the same 5,551 quotes with the same volume- and spread-sensitive price objective. Its coefficients are constant over this one-month window. Both models retain the observed forwards and discount factors; neither is calibrated to forward-start options.

Figure 2 overlays the forward predictions and the observed spot-market smile across strikes. Each

Table 3: At-the-money forward-start comparisons from the October 30, 2026 SPXW snapshot. The premium difference is Heston relative to the clock-based model for the same call payoff. These are model predictions, not traded forward-start quotes.

Reset	Maturity	Clock-based model volatility	Heston volatility	Heston call premium difference
November 2	November 16	16.79%	16.34%	−2.6%
November 9	November 23	15.86%	16.83%	+6.1%
November 16	December 1	14.75%	17.00%	+15.3%

smile uses log-moneyness relative to its own forward; volatility levels are neither shifted nor rescaled, and the market quotes are not interpolated. Both models generate pronounced downside skew and an upturn in the call wing. The distinction is how the forward smile evolves with the reset date: at-the-money forward volatility declines across these examples in the clock-based prototype, whereas it rises slightly under Heston. The call-premium difference changes sign and grows to about 15% at the last reset. For a desk pricing a reset option or a cliquet component, these are different valuations of future risk, even though both descriptions produce familiar smile shapes. A curved call wing by itself is therefore not evidence distinguishing rough dynamics from Heston.

The prototype's forward smiles retain a pronounced downside skew, allowing its shape and strength to be compared directly with the two-week market smile in the same panels. This preservation of the spot smile's character is consistent with the practitioner expectation described above. It is a positive feature for future-reset pricing, in contrast to the flattening found in the published local-volatility comparisons. Heston also preserves substantial skew in this example; the evidence supports this qualitative property of the prototype without establishing a uniform advantage over Heston.

The comparison is to a common market, not to identical fitted marginals. The prototype has the closer vanilla fit and includes a fitted maturity profile; Heston uses five constant coefficients. The difference in forward prices therefore cannot be attributed solely to the rough-clock mechanism.

All three prototype examples use the same 524,288 simulated paths and the numerical grid used for the vanilla illustrations. The at-the-money Monte Carlo standard error is approximately 0.02 volatility percentage points. This quantifies sampling uncertainty; a separate time-step convergence study remains necessary to quantify discretization error. Heston is evaluated for the same cash payoff; doubling its quadrature resolution leaves the displayed figures unchanged.

5.6 Implications for valuation and hedging

The comparison illustrates why a desk needs more than a familiar smile shape. Both models generate skew and curvature, yet disagree on how future-reset exposure evolves across dates. A vanilla fit determines the cost of today's terminal payoffs; the joint dynamics determine the cost of encountering volatility after a future market move. Forward-smile shape, its response to spot and the associated hedge exposure are therefore natural tests of the modelling choice.

Traded reset-option prices or realized hedge results are needed to establish mispricing. The numerical differences reported here quantify model risk; they are not errors measured against observed forward-start quotes.

A complementary test holds structural coefficients fixed across subsequent market dates, updates spot, forwards and time to expiry, and changes only the volatility state. Its ability to explain observed SSR, skew and curvature tests the dynamics beyond a single snapshot. Comparing historical and pricing-measure predictions requires explicit assumptions connecting the two measures.

The contribution is a compact interpretation that generates an equity smile and future-reset prices

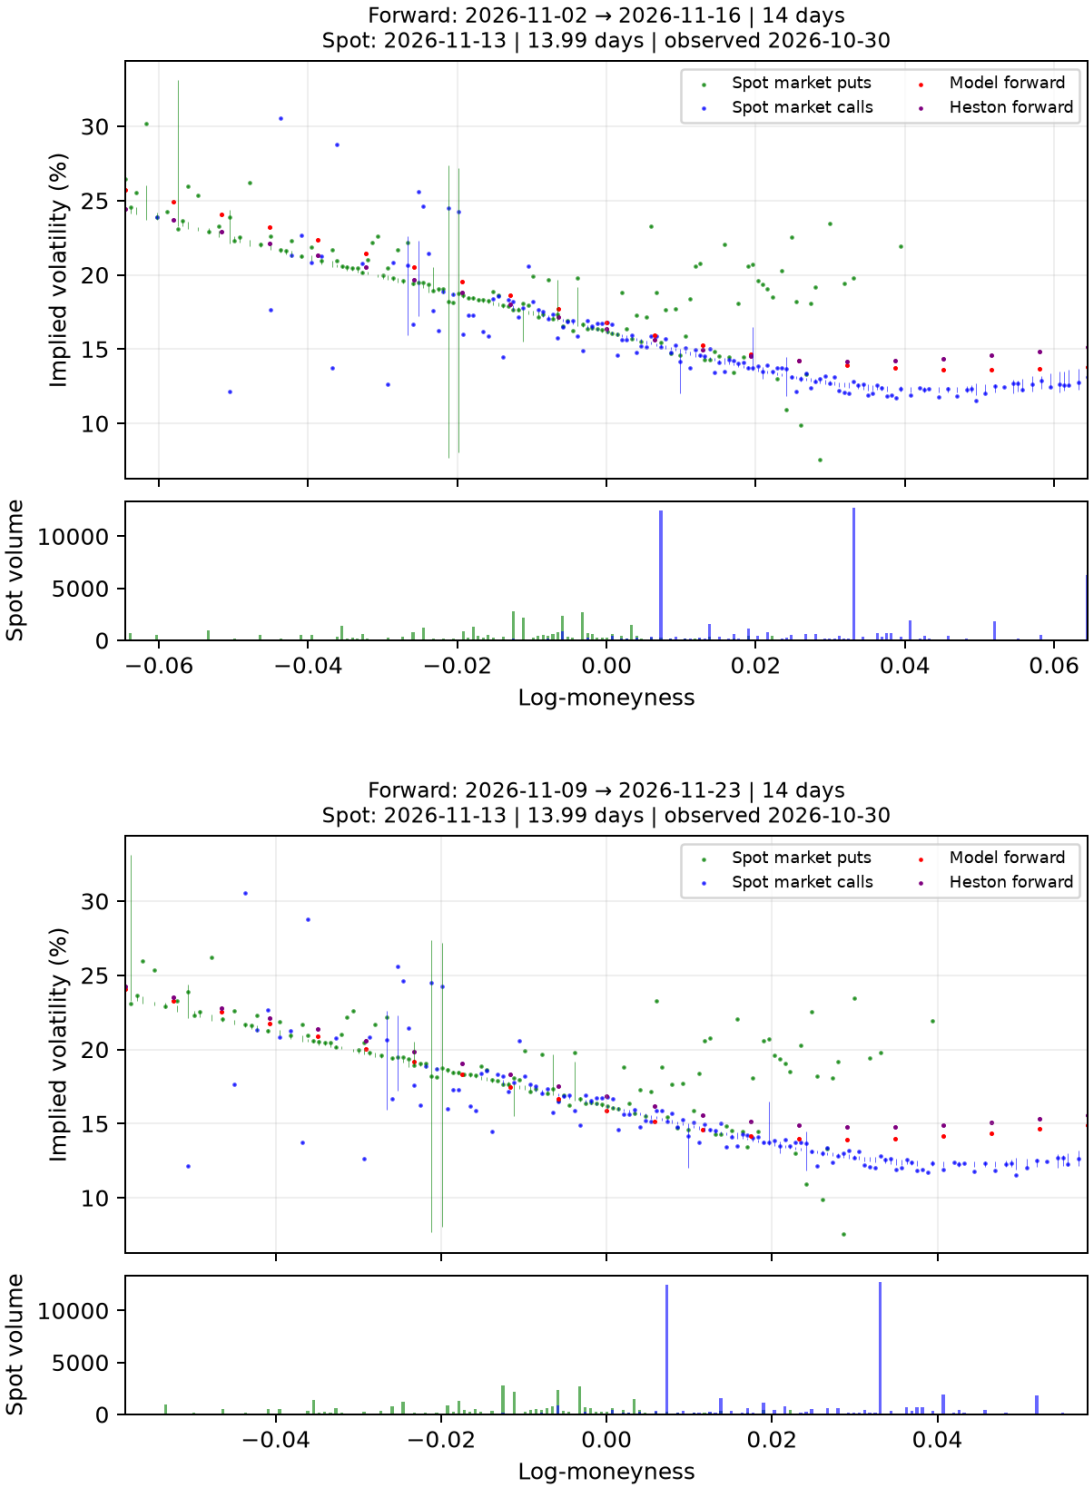


Figure 2: Forward smiles compared directly with the observed spot-market smile. Green and blue points are October 30 closing put and call implied volatilities for the November 13 SPXW expiry (13.99 days); red and purple points are the clock-based and Heston forward-start predictions. The forward option periods in these two panels are 14 days. Market-coloured vertical ranges show bid/ask spreads and lower panels show observed volume; red ranges show one Monte Carlo standard error. Moneyness uses each contract's forward, without shifting volatility levels. The market points are spot-start options, not traded forward-start quotes. Source: ThetaData.

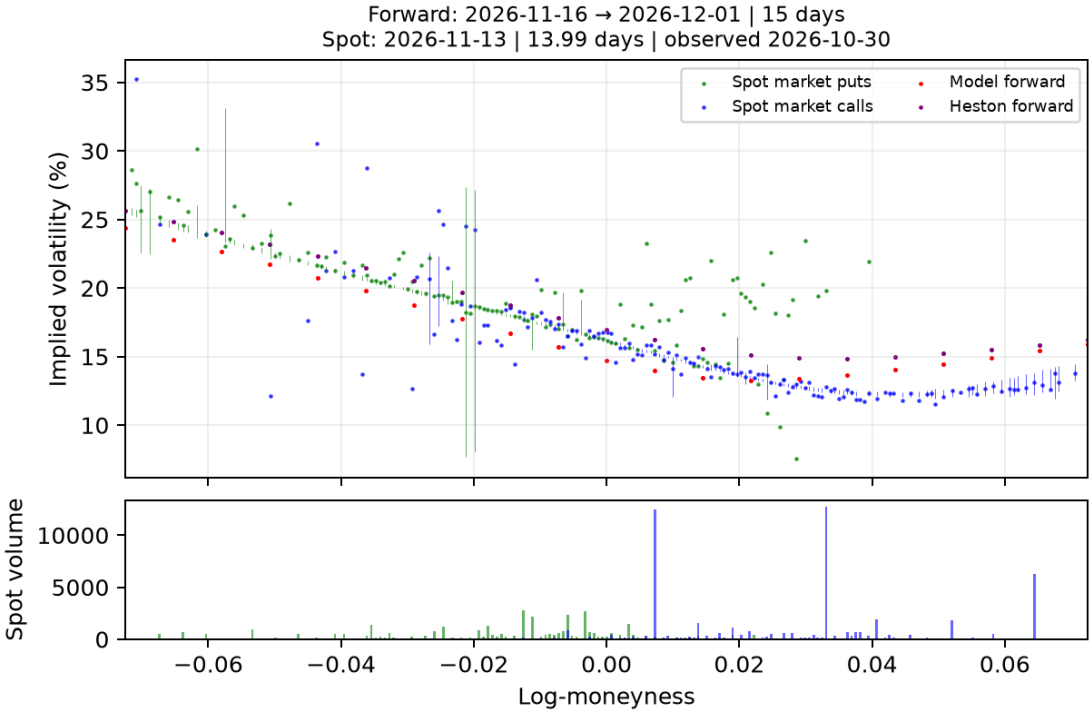


Figure 2: Forward smile versus the spot market, continued: the November 16 reset and December 1 maturity give a 15-day forward option period. The observed spot reference remains the November 13 expiry (13.99 days), the nearest available tenor. Colours and quote conventions are as in the preceding panels. No synthetic 15-day market quotes are introduced.

without a separate state-dependent leverage fit. The SPXW and Heston results demonstrate its practical use and a material difference in predictions. The local- and stochastic-local-volatility results discussed here are from the cited literature; no new numerical LV or SLV comparison is claimed. Empirical tests of smile dynamics and hedging remain the basis for deciding whether the proposed interpretation improves decisions about future-volatility exposure.

References

[1] J. Gatheral, T. Jaisson and M. Rosenbaum (2018). Volatility is rough. *Quantitative Finance*, 18(6), 933–949. arXiv:1410.3394.

[2] C. Bayer, P. Friz and J. Gatheral (2016). Pricing under rough volatility. *Quantitative Finance*, 16(6), 887–904. doi:10.1080/14697688.2015.1099717.

[3] M. Bennedsen, A. Lunde and M. S. Pakkanen (2022). Decoupling the Short- and Long-Term Behavior of Stochastic Volatility. *Journal of Financial Econometrics*, 20(5), 961–1006. doi:10.1093/jjfinec/nbaa049.

[4] P. K. Friz and J. Gatheral (2024). *Computing the SSR*. arXiv:2406.16131.

[5] M. Fukasawa (2021). Volatility has to be rough. *Quantitative Finance*, 21(1), 1–8. arXiv:2002.09215.

[6] L. Bergomi (2009). *Smile Dynamics IV*. doi:10.2139/ssrn.1520443.

[7] L. Bergomi (2005). *Smile Dynamics II*. doi:10.2139/ssrn.1493302.

[8] F. Bourgey, J. Delemotte and S. De Marco (2024; revised 2026). *Smile Dynamics and Rough Volatility*. doi:10.2139/ssrn.4911186.

[9] P. K. Clark (1973). A Subordinated Stochastic Process Model with Finite Variance for Speculative Prices. *Econometrica*, 41(1), 135–155. doi:10.2307/1913889.

[10] G. Pelts (2024). *Roughly Rough: Tractable Markov Processes Emulating Rough Volatility*. doi:10.2139/ssrn.4769827.

[11] G. Pelts (2015). *Liquidity, Trading Volume, and Calibration Target*. ResearchGate preprint. doi:10.13140/RG.2.2.20308.22405.

[12] S. Badikov, A. Jacquier, D. Q. Liu and P. Roome (2016). *No-arbitrage bounds for the forward smile given marginals*. arXiv:1603.06389.

[13] C. Albanese, H. Lo and A. Mijatović (2009). Spectral methods for volatility derivatives. *Quantitative Finance*, 9(6), 663–692. arXiv:0905.2091.

[14] S. L. Heston (1993). A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options. *The Review of Financial Studies*, 6(2), 327–343. doi:10.1093/rfs/6.2.327.

[15] A. W. van der Stoep, L. A. Grzelak and C. W. Oosterlee (2014). The Heston Stochastic-Local Volatility Model: Efficient Monte Carlo Simulation. *International Journal of Theoretical and Applied Finance*, 17(7), 1450045. doi:10.1142/S0219024914500459. Author-hosted full text.