

Quasi-Explicit Recovery of Heston Parameters from SSVI-Parameterized Volatility Surfaces

chris bemis

University of Minnesota, Department of Mathematics

September 14, 2026 — working draft

Abstract

The calibration of the Heston model via least-squares to an implied volatility surface is ill-posed as the surface identifies combinations of the parameters $(\kappa, \theta, \xi, \rho, v_0)$ while failing to disambiguate individual values. We assemble, from published asymptotic results together with several elementary computations, a *quasi-explicit recovery map* with each Heston parameter expressed as an explicit functional of a fitted stochastic volatility inspired surface (SSVI). The construction relies on a sequential and exact identification across two asymptotic regions and the at-the-money term structure, and the identification of each parameter requires no more than scalar monotone root-finding.

1 Introduction

The Heston model [12],

$$dX_t = -\frac{1}{2}v_t dt + \sqrt{v_t} dW_t, \quad (1)$$

$$dv_t = \kappa(\theta - v_t) dt + \xi\sqrt{v_t} dZ_t, \quad d\langle W, Z \rangle_t = \rho dt, \quad (2)$$

provides an options pricing framework with several key features defined specifically by the interpretability of the parameters of the model. These are

- $X_t = \ln(S_t/S_0)$ as the log forward price,
- v_0 as the time zero variance estimate,
- κ as the mean reversion speed of variance and a control term for the term structure of volatility. High κ implies shocks dissipate quickly; low κ implies persistence,
- θ as the long-run variance,
- ξ as the volatility of volatility, controlling the kurtosis (or tail risk) of the asset distribution,
- ρ as the correlation between changes in the underlying and changes in variance, typically negative, resulting in rising variance with decreasing underlying, and controlling the skew of the options curve.

However, its calibration is governed by an objective function that is flat along several parameter combinations, making any surface least squares formulation correspondingly ill-posed. As a result, optimizers deliver optimal solutions pinned at boundary conditions or pairwise combinations that minimize the objective but lose interpretability. For example, as is widely documented in the quantitative finance literature [3, 8, 17], the market smile robustly identifies combinations such as the skew-governing product $\rho\xi$, but leaves the individual specification of ρ and ξ unsettled. Even when an identification may be possible, numerical routines frequently place ρ at arbitrary bounds.

The SSVI surface of Gatheral and Jacquier [10] presents the opposite profile, with three global shape parameters (ρ_s, η, γ) fit by a well-behaved low-dimensional problem via quasi-explicit no-arbitrage conditions, yielding stable, initial point invariant estimates. As a result, the shape parameters satisfy both uniqueness as well as *a priori* interpretability. The present note takes SSVI as the measurement layer and asks for the complete map

$$\text{SSVI fit } (\rho_s, \eta, \gamma; \{\vartheta_T\}) \longmapsto (\kappa, \theta, \xi, \rho, v_0),$$

with $\{\vartheta_T\}$ a particular implied variance curve indexed by maturity, T . The map is motivated by a fact that elevates SSVI beyond just a smoothing convenience. Namely, that *the large-maturity limit of the Heston smile is exactly SVI* [9], and the small-time Heston smile is likewise explicit [5, 7]. SSVI therefore belongs to Heston's own family in the relevant asymptotic regimes. Reading Heston parameters off an SSVI fit is therefore not a heuristic bridge between unrelated objects, but an inversion of published asymptotics evaluated on the parameterizations to which they converge.

This inversion constitutes our *recovery map*, directly obtaining stable and individually interpretable Heston components from the well-behaved and optimal SSVI parameters. The recovery map is sequential and deterministic, and nowhere relies on a multidimensional search. In what follows, Sections 3–5 develop the three regimes we leverage to isolate particular features of the Heston model. In doing so, and wherever possible without assumption, we show that they are limits of one *analytic* object rather than three separate approximations. Section 6 formalizes this translation into a step-by-step mapping, proves its point-identifiability, and discusses the model’s degenerate limits; e.g., considers the reduction to a Black–Scholes assumption and exhibits model robustness.

Section 7 turns to implementation. A specific contrast between the method developed herein and an optimization procedure is that the various regimes utilized provided an over-identification of some key parameters. Measured gaps between such over-identified parameters may be utilized as specification statistics, replacing regularization. In addition, we identify in this section the bias of using a finite set of tenors and the impact of choosing a level curve on a fixed-delta slice.

Section 8 tests the map on surfaces of known parameters in the protocol of the quasi-explicit SVI calibration [16] from which we take the term. Here we also exhibit that the calibration of the mean reversion parameter, κ , cannot be obtained from a global power law assumption. Rather, the skew term structure must be read slice by slice, otherwise the result is a read on the exponent γ and not κ . This potentially impacts the literature on the ubiquity of the so-called Feller violation, which we discuss further below.

Section 9 concludes.

2 Preliminaries

2.1 Model and conventions

We focus on the Heston model outlined by the system (1)–(2) with $\kappa, \theta, \xi, v_0 > 0$, $\rho \in (-1, 1)$. We do not enforce or assume the Feller condition, $2\kappa\theta \geq \xi^2$, which is often required for the stability of standard Euler discretization schemes [1, 15]. We will, however, note a specific impact of our work in the measurement of market data violations of this condition in Section 8. Elsewhere, the Feller condition does not impact any read from any regime, asymptotic or otherwise.

For log-moneyness k and maturity T , let $\sigma(k, T)$ denote Black–Scholes implied volatility and define the total implied variance as

$$w(k, T) = \sigma^2(k, T) T. \quad (3)$$

The at-the-money implied variance is obtained by setting $k = 0$, and we define the at-the-money implied variance skew as

$$\varsigma(T) := \partial_k (\sigma^2(k, T)) \big|_{k=0}. \quad (4)$$

2.2 The SSVI family

The SSVI surface introduced by Gatheral and Jacquier [10] parameterizes each time slice by its at-the-money total implied variance

$$\vartheta_T := w(0, T).$$

For shape parameters $\rho_s \in (-1, 1)$ and $\varphi > 0$, the SSVI formulation is given by

$$w_{SSVI}(k, \vartheta_T) = \frac{\vartheta_T}{2} \left(1 + \rho_s \varphi k + \sqrt{(\varphi k + \rho_s)^2 + 1 - \rho_s^2} \right), \quad (5)$$

with $\varphi = \varphi(\vartheta_T)$. To guarantee the absence of static arbitrage and capture empirical skew features, we adopt the canonical *power-law specification*

$$\varphi(\vartheta) = \eta \vartheta^{-\gamma}$$

advocated in [10].

As established below, the function φ provides the analytic bridge to the Heston skew, embedding the volatility-of-volatility ξ and mean reversion κ into the term structure. We have an explicit connection given by the following evaluations of (5),

$$\partial_k w_{SSVI}|_{k=0} = \rho_s \varphi(\vartheta_T) \vartheta_T, \quad w_{SSVI}(k, \vartheta_T) \sim \frac{\vartheta_T \varphi(\vartheta_T)(1 \pm \rho_s)}{2} |k| \quad (k \rightarrow \pm\infty). \quad (6)$$

Following Gatheral and Jacquier [10], the first expression formalizes the at-the-money implied variance skew, while the second characterizes the asymptotic left and right wing slopes in the extreme-strike limits. These will be utilized in the construction that follows.

2.3 The architectural roadmap

In contrast to the SSVI model established above, Heston surface optimization determined over a calibration window (say, $[0, T_{\max}]$ with $\kappa T_{\max} \lesssim 1$) yields a notoriously ill-posed objective function that exposes blended parameter projections rather than isolated components. Specifically, the variance expectation $\mathbb{E}[v_t] = \theta + (v_0 - \theta)e^{-\kappa t}$ forces the (κ, θ) pair to slide freely along a ridge of constant variance level. Similarly, the smile slope locks the product $\rho\xi$ at leading order, but fails to differentiate the leverage correlation from the volatility-of-volatility. Any numerical optimizer delivering a single Heston point estimate in any of these cases without further specification or constraint is introducing unspecified arbitrariness.

Perhaps counterintuitively, the remedy is to abandon direct joint surface optimization entirely. By using the stable SSVI surface as our measurement layer, we may break the Heston structural dependencies by reading each parameter exclusively from the asymptotic regime where its unique identification is proved. The remainder of this paper executes this decoupling to identify the full Heston parameter set, $(\kappa, \theta, \xi, \rho, v_0)$, as outlined below.

- **The small-time regime:** Isolates the spot variance v_0 and the *combined* skew product $\rho\xi$ at the short end, $T \rightarrow 0$. The *level* determines v_0 and the slope $\rho\xi$.
- **At-the-money term structures:** Breaks the (κ, θ) ridge by identifying the mean reversion κ strictly through the *shape* of the skew decay, separating it entirely from the variance level. The *shape* of the skew curve isolates κ and, once fixed, the *asymptote* gives θ .
- **The large-maturity regime:** Leverages the exact SVI limit at $T \rightarrow \infty$ to isolate the correlation ρ , allowing a factorization of the $\rho\xi$ product obtained at the short end and thus giving a full specification of the Heston model. The *asymmetry* of the wings gives ρ .

Section 6 formalizes the quasi-explicit recovery map, replacing ad hoc parameter fixing with continuous specification testing utilizing this sequential mapping.

3 The small-time regime

Consider the local variance function

$$\sigma_{loc}^2(k, T) = \mathbb{E}[v_T \mid X_T = k]. \quad (7)$$

Proposition 3.1. *To first order in log-moneyness k and small maturity T , the implied variance expands as*

$$\sigma^2(k, T) = v_0 + \frac{1}{2}\rho\xi k + O(T) + O(k^2).$$

Proof. For small T , v_T is tightly concentrated near v_0 , effectively freezing the diffusion coefficients. Consequently, the local increments of the process pair (X_T, v_T) are dominated by the correlated Brownian motions (W_T, Z_T) and thereby admit a joint Gaussian approximation.

Applying the Itô isometry directly to the system (1)–(2) under this constant-coefficient limit, the leading-order variance and covariance are $\text{Var}(X_T) = v_0 T + O(T^2)$ and $\text{Cov}(X_T, v_T) = \rho\xi v_0 T + O(T^2)$. Linear projection of v_T onto X_T then gives the conditional expectation as

$$\sigma_{loc}^2(k, T) = v_0 + \frac{\rho\xi v_0 T}{v_0 T} k + O(T) + O(k^2) = v_0 + \rho\xi k + O(T) + O(k^2).$$

By the Berestycki–Busca–Florent expansion [2], the small-time implied volatility is the harmonic mean of the local volatility between the spot and the strike. At first order in k every such average of a function with slope $\rho\xi$ at the origin has slope $\frac{1}{2}\rho\xi$, so

$$\sigma^2(k, T) = v_0 + \frac{1}{2}\rho\xi k + O(T) + O(k^2),$$

concluding the proof. $\square$

Proposition 3.2 (Level). $\lim_{T \rightarrow 0} \sigma^2(k, T) = v_0$ uniformly for k in a neighbourhood of 0. In particular, $\sigma^2(0, T) = v_0 + O(T)$.

Proof. Evaluating the expansion in Proposition 3.1 at the money ($k = 0$) isolates the spot variance and preserves the time correction, yielding $\sigma^2(0, T) = v_0 + O(T)$. The rigorous bounds for this $O(T)$ departure are structurally confirmed by the large-deviations rate function in Forde and Jacquier [5]. $\square$

Proposition 3.3 (Skew). $\lim_{T \rightarrow 0} \varsigma(T) = \frac{1}{2}\rho\xi$.

Proof. Differentiating the expansion in Proposition 3.1 with respect to k and evaluating at $k = 0$ yields $\varsigma(T) = \frac{1}{2}\rho\xi + O(T)$. Taking the limit as $T \rightarrow 0$ proves the claim. $\square$

Proposition 3.4 (Curvature). $\partial_k^2 \sigma^2(k, T)|_{k=0}$ converges as $T \rightarrow 0$ to an explicit constant of the form $c(\rho)\xi^2/v_0$, with $c(\cdot)$ a known rational function of ρ given by Forde–Jacquier–Lee [7].

Proof. The joint Gaussian approximation of Proposition 3.1 yields a linear conditional expectation and is therefore insufficient at second order in k . Capturing the k^2 dependence requires the second-order conditional moment, which is governed by the variance of the variance process itself. From the Itô dynamics, $\text{Var}(v_T) = \xi^2 v_0 T + O(T^2)$. Projecting this non-Gaussian quadratic variation onto the log-moneyness X_T (which has variance $v_0 T$) introduces a curvature term that must scale dimensionally as ξ^2/v_0 , modulated by the correlation ρ .

This dimensional scaling establishes the proposition’s claim, leaving only the exact algebraic form of the correlation multiplier $c(\rho)$ to be evaluated, which depends on the second-order form of the small-time expansion and is done explicitly in Forde, Jacquier, and Lee [7]. $\square$

We identify Proposition 3.4 here not because it is required for the core parameter recovery. In fact, the remaining parameters are fully identified by the term structure and long-maturity regimes. We include the result as it provides a built-in, observable test for model misspecification detailed in §7.1, where short-end curvature is compared against the volatility-of-volatility ξ recovered from the broader surface.

4 At-the-money term structures

With the spot variance v_0 and the combined skew product $\rho\xi$ anchored by the short-time regime, the recovery map next extracts the mean reversion κ and the long-run variance θ . A natural (though ultimately insufficient) starting point is the term structure of the at-the-money variance level.

To leading order, the ATM total implied variance ϑ_T is governed by the expected integrated variance of the process. We formalize this baseline expectation next in terms of the Heston parameters for long term variance and reversion (θ and κ , respectively) given in (1)–(2).

Lemma 4.1. $\mathbb{E}[v_t] = \theta + (v_0 - \theta)e^{-\kappa t}$, and consequently the expected total integrated variance is

$$\mathcal{W}(T) := \mathbb{E} \int_0^T v_t dt = v_0 B(\kappa, T) + \theta(T - B(\kappa, T)), \quad B(\kappa, T) = \frac{1 - e^{-\kappa T}}{\kappa}. \quad (8)$$

Proof. Recall the stochastic differential equation governing the variance process from (2):

$$dv_t = \kappa(\theta - v_t) dt + \xi \sqrt{v_t} dZ_t$$

Integrating this from 0 to t yields the integral equation:

$$v_t = v_0 + \int_0^t \kappa(\theta - v_s) ds + \int_0^t \xi \sqrt{v_s} dZ_s$$

Taking expectations on both sides, we first address the stochastic integral. Because $\mathbb{E} \int_0^t v_s ds < \infty$ (which can be formally established via comparison with the linear deterministic SDE), the stochastic integral $\int_0^t \xi \sqrt{v_s} dZ_s$ is a true martingale. Its expectation therefore vanishes, leaving

$$\mathbb{E}[v_t] = v_0 + \kappa\theta t - \kappa \int_0^t \mathbb{E}[v_s] ds$$

where Fubini's theorem allows the interchange of the expectation and the Riemann integral.

Defining the deterministic function $m(t) = \mathbb{E}[v_t]$, we differentiate with respect to t to obtain a linear ordinary differential equation

$$m'(t) = \kappa(\theta - m(t))$$

subject to the initial condition $m(0) = v_0$. Multiplying by the integrating factor $e^{\kappa t}$ leads to the solution,

$$m(t) = \theta + (v_0 - \theta)e^{-\kappa t}.$$

To find the expected total integrated variance $\mathcal{W}(T)$, we again invoke Fubini's theorem (justified by the non-negativity of the variance process) to bring the expectation inside the integral and get

$$\mathcal{W}(T) = \mathbb{E} \int_0^T v_t dt = \int_0^T \mathbb{E}[v_t] dt = \int_0^T (\theta + (v_0 - \theta)e^{-\kappa t}) dt.$$

Evaluating the integral directly gives the claimed result, completing the proof. $\square$

While (8) provides a strict structural relationship, it also explicitly manifests the calibration ridge previewed in Section 2.3. Equating an observed variance curve ϑ_T with $\mathcal{W}(T)$ on a finite window constrains (κ, θ) to a nearly one-dimensional manifold of constant variance level, and cannot point-identify either parameter in isolation with any useful conditioning (Remark 4.3). This degeneracy may be resolved, however, if we examine the maturity decay of the variance *skew* and not the variance *level*. We arrive at a theorem for the at-the-money skew term structure after first proving the following proposition.

Proposition 4.2. *The at-the-money total variance, $w(0, T)$, satisfies $w(0, T) = \mathcal{W}(T) (1 + O(\xi^2))$, uniformly on compact T -sets. In particular the map $T \mapsto w(0, T)$ on $(0, \infty)$ identifies (v_0, θ, κ) through (8) up to $O(\xi^2)$.*

Proof. We begin by expanding the implied total variance $w(k, T)$ in L^2 for small volatility-of-volatility ξ . As established by Lewis [14] and Gatheral [8], this expansion takes the form

$$w(k, T) = w^{(0)}(k, T) + \xi w^{(1)}(k, T) + O(\xi^2).$$

At the zeroth order ($\xi = 0$), the Heston stochastic differential equation collapses to the deterministic ordinary differential equation $dv_t^{(0)} = \kappa(\theta - v_t^{(0)}) dt$ solved in Lemma 4.1. Consequently, the implied total variance reduces trivially to the deterministic integrated variance, giving $w^{(0)}(k, T) = \mathcal{W}(T)$.

The perturbation $w^{(1)}(k, T)$ is driven entirely by the spot-variance correlation and is strictly linear in the log-moneyness k . Since this term is an *odd* function of k , it vanishes identically at the money ($k = 0$), yielding $w^{(1)}(0, T) = 0$. Substituting these into the expansion directly provides our first claim,

$$w(0, T) = \mathcal{W}(T) + O(\xi^2) = \mathcal{W}(T)(1 + O(\xi^2)).$$

To demonstrate how this isolates the parameter subset as claimed, we recall from Lemma 4.1 that

$$\mathcal{W}(T) = \theta T + (v_0 - \theta) \frac{1 - e^{-\kappa T}}{\kappa}. \quad (9)$$

On the full half-line the mapping $T \mapsto \mathcal{W}(T)$ uniquely characterizes the drift parameters through its geometric properties across the term structure. The short-term behavior is anchored by the spot variance $(d\mathcal{W}(T)/dT|_{T=0} = v_0)$, the asymptotic long-term slope gives the mean reversion level $(\lim_{T \rightarrow \infty} \mathcal{W}(T)/T = \theta)$, and the exponential curvature defines the mean reversion rate κ . Because these three features are independent, observing the at-the-money variance $w(0, T)$ across all maturities uniquely identifies the triplet (v_0, θ, κ) up to the structural $O(\xi^2)$ perturbation, thus completing the proof. $\square$

Remark 4.3. Proposition 4.2 identifies the triplet (v_0, θ, κ) given the variance level curve, but assumes the infinite horizon $(0, \infty)$. In practice, observable data is restricted to a finite tenor window $[0, T_{\max}]$, however. Here, the structural identification deteriorates into the calibration ridge discussed in Section 2.3.

To see why, consider the geometric features captured within $[0, T_{\max}]$. The short end anchors the spot variance via the derivative $v_0 = \mathcal{W}'(0)$ and the long end provides the average

variance out to the horizon, $\bar{v} := \mathcal{W}(T_{\max})/T_{\max}$. The transition between these boundaries yields an initial curvature proportional to the product $\kappa(\theta - v_0)$. The true asymptote $\theta = \lim_{T \rightarrow \infty} \mathcal{W}(T)/T$ lies *strictly outside the observable data*, however. If the mean reversion is slow relative to the window ($\kappa T_{\max} \lesssim 1$), the decay has barely begun by $T_{\max}$ and the asymptote is purely an extrapolation.

Attempting to extract θ and κ simultaneously from any finite strip, therefore, equates to fitting an exponential asymptote to data that truncates before the decay has run its course. This is a classically ill-conditioned inverse problem. While the parameters remain theoretically unique, the level curve resolves (κ, θ) with numerical fragility. It is this instability that dictates our architectural sequence; viz., the recovery map must extract κ from the term structure of the skew, reserving the variance level curve strictly to project θ *conditioned upon a κ that is already firmly anchored*.

The variance term structure constrains the drift parameters, but on a finite window it cannot resolve the degeneracy between the mean reversion κ and the long-run variance θ with useful conditioning. To break this degeneracy, we must look beyond the variance level and examine its dynamics across maturities of the at-the-money implied variance skew. Our immediate objective is not to decouple the correlation ρ and the volatility-of-volatility ξ , which remain bounded here as a single scaling coefficient, but to demonstrate that the temporal decay of the skew isolates κ entirely through its shape.

Assumption 1. *In the first-order approximation of the variance skew, we assume the expected variance $\mathbb{E}[v_s]$ governing the covariance $\text{Cov}(v_t, X_T)$ may be replaced by a constant effective variance $\bar{v}$, which identically scales the terminal variance $\text{Var}(X_T)$. This substitution is exact for a flat forward variance curve ($v_0 = \theta$) and holds to leading order in the relative gap $(v_0 - \theta)/\theta$ otherwise.*

Theorem 4.4. *Under Assumption 1, at first order in ξ , the at-the-money variance skew is given by*

$$\varsigma(T) = \rho \xi g(\kappa T) + O(\xi^2), \quad (10)$$

where

$$g(x) = \frac{1}{x} \left(1 - \frac{1 - e^{-x}}{x} \right) \quad (11)$$

with $g(0^+) = \frac{1}{2}$ and $g(x) \sim \frac{1}{x}$ ($x \rightarrow \infty$).

Proof. At first order in the volatility-of-volatility ξ , the implied variance admits the conditional-mean representation established by the Bergomi–Guyon expansion [3]. Specifically, it equates to the time-averaged expected variance conditioned on the terminal state as

$$\frac{w(k, T)}{T} = \frac{1}{T} \int_0^T \mathbb{E}[v_t | X_T = k] dt + O(\xi^2).$$

To evaluate this conditional expectation, we note that to leading order, the intermediate variance v_t and the terminal log-return X_T are jointly Gaussian. Gaussian conditioning therefore allows us to express the expectation as a linear function of the log-moneyness k ,

$$\mathbb{E}[v_t | X_T = k] = \mathbb{E}[v_t] + \frac{\text{Cov}(v_t, X_T)}{\text{Var}(X_T)} k + O(\xi^2).$$

This structural relation reduces the problem to computing the covariance between the variance path and the terminal return.

We apply the variation of constants to the variance SDE (2) to isolate the stochastic integral as

$$v_t = \mathbb{E}[v_t] + \xi \int_0^t e^{-\kappa(t-s)} \sqrt{v_s} dZ_s.$$

To leading order, the log-return is driven by its continuous martingale part, $X_T \approx \int_0^T \sqrt{v_u} dW_u$. Applying the Itô isometry with the correlation structure $d\langle Z, W \rangle_t = \rho dt$, and recognizing that increments of X after time t are independent of v_t at this order, we obtain

$$\text{Cov}(v_t, X_T) = \mathbb{E} \left[\left(\xi \int_0^t e^{-\kappa(t-s)} \sqrt{v_s} dZ_s \right) \left(\int_0^T \sqrt{v_u} dW_u \right) \right] = \rho \xi \int_0^t e^{-\kappa(t-s)} \mathbb{E}[v_s] ds + O(\xi^2).$$

Under Assumption 1, replacing $\mathbb{E}[v_s]$ by the effective level $\bar{v}$, this integral evaluates to

$$\text{Cov}(v_t, X_T) = \rho \xi \bar{v} B(\kappa, t) + O(\xi^2).$$

With the terminal variance given by $\text{Var}(X_T) = \bar{v}T + O(\xi)$ under the same assumption, we may substitute these moments back into the Gaussian conditioning formula and determine that

$$\mathbb{E}[v_t | X_T = k] = \mathbb{E}[v_t] + \frac{\rho \xi B(\kappa, t)}{T} k + O(\xi^2),$$

where the effective variance level $\bar{v}$ cancels in the slope term.

Substituting this conditional expectation back into the Bergomi–Guyon representation yields the annualized implied variance as a direct function of the log-moneyness k ,

$$\frac{w(k, T)}{T} = \frac{1}{T} \int_0^T \left(\mathbb{E}[v_t] + \frac{\rho \xi B(\kappa, t)}{T} k \right) dt + O(\xi^2).$$

The ATM variance skew $\varsigma(T)$ is defined as the log-moneyness derivative of this annualized implied variance evaluated at $k = 0$. Differentiating the assembled expression with respect to k eliminates the zeroth-order expectation and isolates the integral of the slope. We have

$$\varsigma(T) = \frac{\partial}{\partial k} \left(\frac{w(k, T)}{T} \right) \Big|_{k=0} = \frac{1}{T} \int_0^T \frac{\rho \xi B(\kappa, t)}{T} dt + O(\xi^2).$$

Evaluating the deterministic integral $\int_0^T B(\kappa, t) dt = (T - B(\kappa, T))/\kappa$ gives the result of the proposition as

$$\varsigma(T) = \frac{\rho \xi}{T^2} \cdot \frac{T - B(\kappa, T)}{\kappa} + O(\xi^2) = \rho \xi g(\kappa T) + O(\xi^2). \quad \square$$

If Assumption 1 is relaxed, propagating the exact variance term structure $m(s) = \mathbb{E}[v_s]$ through the covariance integral yields the generalized first-order skew

$$\varsigma(T) = \rho \xi \frac{\int_0^T \int_0^t e^{-\kappa(t-s)} m(s) ds dt}{T \int_0^T m(u) du} + O(\xi^2). \quad (12)$$

This trivially reduces to (10) when the variance level is constant. When the forward curve is not flat, however, the relative gap $(v_0 - \theta)/\theta$ alters the skew decay profile. Consequently, extracting κ via the simplified shape function $g(\kappa T)$ introduces a bias proportional to this gap.

Fortunately, the sequential architecture of our recovery map accommodates this natively. To address the bias, the map can simply evaluate (12) using the (v_0, θ) values obtained from the baseline recovery, updating κ through a single fixed-point iteration. This scalar correction de-biases the mean reversion without a multidimensional search.

Of course, the results of Theorem 4.4 obtain at first order and the proof explicitly holds at that resolution. Higher-order expansions [8] demonstrate that the empirical skew decay is governed by an effective mean reversion rate, $\kappa_{\text{eff}} = \kappa - \rho\xi/2$. Because our sequential recovery map already anchors the combined product $\rho\xi$ from the short-time regime, this second-order bias is completely observable. Consequently, the prescription that follows for the identification of Heston parameters in our recovery map may natively incorporate this $g(\kappa_{\text{eff}}T)$ formulation.

Finally, we show that the skew term structure identifies κ *alone*, breaking a degeneracy that the level curve inherently cannot.

Lemma 4.5. *Suppose $a_1 g(\kappa_1 T) = a_2 g(\kappa_2 T)$ for all T in some open interval (T_1, T_2) , with $a_i \neq 0$ and $\kappa_i > 0$. Then $\kappa_1 = \kappa_2$ and $a_1 = a_2$.*

Proof. The function g extends to an entire function $g(x) = \sum_{n=0}^{\infty} c_n (-x)^n$ with $c_n = \frac{1}{(n+2)!} > 0$ via the series expansion $g(x) = \frac{1}{2} - \frac{x}{6} + \frac{x^2}{24} - \dots$. Because $T \mapsto a g(\kappa T)$ is real-analytic, agreement on an open interval forces all Taylor coefficients to match: $a_1 c_n \kappa_1^n = a_2 c_n \kappa_2^n$ for all $n \geq 0$. Equating the $n = 0$ coefficients yields $a_1 = a_2$. Since $a_1 \neq 0$, equating the $n = 1$ coefficients then forces $\kappa_1 = \kappa_2$. $\square$

As a result, the *shape* of the observed skew curve isolates κ , while its *scale* determines the product $\rho\xi$. While we have not as of yet disambiguated ρ and ξ , the preceding analysis does identify κ outside of the degeneracy of the level curve of (κ, θ) . The mean reversion, κ , so identified is expressly determined by *skew persistence* and not *variance level*.

Notice in the work above that at $T \rightarrow 0^+$, we have $g(0^+) = \frac{1}{2}$ and recover Proposition 3.3 exactly, demonstrating that the small-time and term-structure regimes agree where they meet. At large T , the function yields $\varsigma(T) \sim \rho\xi/(\kappa T)$, recovering the classical $1/T$ Heston long-maturity skew decay documented by Gatheral [8] and formalized by Forde and Jacquier [5].

The agreement between the small-time limit and the term-structure derivation is the theoretical foundation of the recovery map identified in this work. The three regimes outlined in Section 2.3 do not represent three distinct approximations of the volatility surface with unique parameters. Rather, they represent the boundary and interior evaluations of a single, unified function $T \mapsto \rho\xi g(\kappa T)$. Because this function is entire in T (as established in Lemma 4.5), its domain connects the calibration regimes. Its evaluation at $T \rightarrow 0^+$ yields the small-time skew of Proposition 3.3, its $1/T$ tail governs the large-maturity decay, and its interior shape completely isolates κ .

Strict analyticity gives license to our methodology as an entire function is determined by its values on any open interval. Because of this, the determination of κ from a finite window of liquid tenors yields the same κ governing the small-time and large-maturity limits. The same principle holds for the variance level curve (8), which is also entire in T . Here, the level curve is anchored by v_0 at the origin, θ at the asymptote, and driven by κ in its interior. The

sequential recovery map does not merely read parameters where they are uniquely defined, but where they are *best conditioned*, leveraging the analytic continuity of a single structural object.

5 The large-maturity regime and the wings

With the drift parameters (v_0, θ, κ) structurally identified and the scaling product $\rho\xi$ anchored by the short-time skew, the recovery map is complete upon disambiguating the correlation ρ from the volatility-of-volatility ξ . The previous section demonstrated that the at-the-money term structures are insufficient to decouple this product.

In the present section, we turn to the asymptotic behavior of the volatility smile at large maturities ($T \rightarrow \infty$) and extreme strikes (the wings). In this large-maturity regime, the Heston implied variance surface converges to a stationary spatial shape, and the asymptotic cumulant generating function (CGF) of the log-return will govern its determination.

We begin by formalizing the limiting CGF.

5.1 The limiting cumulant generating function

Proposition 5.1. *Let $\Lambda(p)$ denote the limiting cumulant generating function of the annualized log-return, defined as $\Lambda(p) = \lim_{T \rightarrow \infty} \frac{1}{T} \ln \mathbb{E}[e^{pX_T}]$.*

For any moment p strictly within the critical interval (p_-, p_+) where the discriminant

$$d(p) := \sqrt{(\kappa - \rho\xi p)^2 + \xi^2(p - p^2)}$$

is real and positive, this asymptotic growth rate converges to the closed form

$$\Lambda(p) = \frac{\kappa\theta}{\xi^2} \left(\kappa - \rho\xi p - d(p) \right). \quad (13)$$

Proof. Since the Heston model belongs to the affine class, the moment generating function admits the exponentially affine form

$$\mathbb{E}e^{pX_T} = \exp(\phi(T) + \psi(T)v_0).$$

The deterministic functions $\phi(T)$ and $\psi(T)$ are characterized by a system of ordinary differential equations subject to the initial conditions $\phi(0) = 0$ and $\psi(0) = 0$ given by

$$\begin{aligned} \psi'(T) &= \frac{1}{2}(p^2 - p) - (\kappa - \rho\xi p)\psi(T) + \frac{\xi^2}{2}\psi(T)^2 \\ \phi'(T) &= \kappa\theta\psi(T). \end{aligned} \quad (14)$$

The equation for ψ is a standard Riccati ODE (see, for instance, Duffie, Pan, and Singleton [4]). Setting the right-hand side of (14) to zero, the quadratic formula yields two equilibrium roots:

$$\psi_{\pm} = \frac{(\kappa - \rho\xi p) \pm \sqrt{(\kappa - \rho\xi p)^2 + \xi^2(p - p^2)}}{\xi^2} = \frac{(\kappa - \rho\xi p) \pm d(p)}{\xi^2}.$$

For p strictly within the critical interval (p_-, p_+) , the discriminant $d(p)$ is strictly positive, ensuring two distinct real roots $\psi_- < \psi_+$.

Since the coefficient for the ψ^2 term in (14) is strictly positive, the quadratic opens upward. Consequently, the vector field ψ' is positive for $\psi < \psi_-$ and negative for $\psi_- < \psi < \psi_+$. The

lower root ψ_- therefore acts as a stable attractor. Starting from the initial condition $\psi(0) = 0$, the trajectory $\psi(T)$ remains bounded and monotonically approaches the stable equilibrium, giving $\lim_{T \rightarrow \infty} \psi(T) = \psi_-$.

Finally, we may evaluate the limiting cumulant generating function. Taking the logarithm of the affine form and dividing by T yields

$$\frac{1}{T} \ln \mathbb{E} e^{pX_T} = \frac{\phi(T)}{T} + \frac{\psi(T)v_0}{T}.$$

As $T \rightarrow \infty$, the fact that $\psi(T)$ is bounded forces the second term, $\frac{\psi(T)v_0}{T}$, to vanish. This structurally confirms that the long-maturity CGF is completely independent of the initial spot variance, v_0 .

To evaluate the remaining ϕ term, we observe that its derivative $\phi'(T) = \kappa\theta\psi(T)$ converges to a finite constant $\kappa\theta\psi_-$ as $T \rightarrow \infty$. Because this limit exists, we may apply L'Hôpital's rule to the ratio $\phi(T)/T$, giving a direct step-by-step evaluation of the limit:

$$\begin{aligned} \Lambda(p) &= \lim_{T \rightarrow \infty} \left(\frac{\phi(T)}{T} \right) \\ &= \lim_{T \rightarrow \infty} \phi'(T) \\ &= \lim_{T \rightarrow \infty} \kappa\theta\psi(T) \\ &= \kappa\theta\psi_-. \end{aligned}$$

Substituting the explicit formula for the stable root ψ_- produces the right side of (13), completing the proof. $\square$

5.2 Convergence of the smile to SVI

With the asymptotic cumulant generating function formally established, we now possess the necessary input to determine the spatial shape of the long-maturity smile. The mathematical bridge between the CGF and the implied variance surface is provided by large deviations theory. Specifically, the Gärtner–Ellis theorem applies a Legendre transform to $\Lambda(p)$ to determine the asymptotic decay rate of the return distribution. By matching these asymptotic option prices to the Black–Scholes formula, Gatheral and Jacquier [9] demonstrated that the limiting Heston smile resolves into a closed-form parametric slice. Because the implied variance at any fixed k asymptotically flattens to the at-the-money level, the smile's shape is preserved only when strikes are scaled proportionally with maturity. As such, this limiting parametric shape emerges under the rescaled log-moneyness coordinate $x = k/T$.

Theorem 5.2 (Gatheral–Jacquier [9]). *Let $x = k/T$ denote the log-moneyness scaled by maturity. As $T \rightarrow \infty$, for each fixed x , the annualized implied variance converges,*

$$\frac{w(xT, T)}{T} \longrightarrow \frac{\omega_1}{2} \left(1 + \omega_2 \rho x + \sqrt{(\omega_2 x + \rho)^2 + 1 - \rho^2} \right). \quad (15)$$

This limit is exactly an SVI slice in x whose correlation parameter equals the Heston ρ , and whose scale parameters are determined by:

$$\omega_1 = \frac{4\kappa\theta}{\xi^2(1 - \rho^2)}, \quad \omega_2 = \frac{\xi}{\kappa\theta}. \quad (16)$$

Summary of proof. By the Gärtner–Ellis theorem applied to the cumulant generating function in Proposition 5.1, the scaled return X_T/T satisfies a large deviations principle with rate function $\Lambda^*(x) = \sup_p (px - \Lambda(p))$ (see Forde and Jacquier [6]). The large-maturity implied variance is an explicit transform of this rate function Λ^* .

Gatheral and Jacquier [9] carry out this Legendre computation in closed form. The substantive algebra of their proof verifies that the resulting function of x matches the SVI parametric form with the scale parameters given in (16). We refer to [9] for the algebraic derivation, but we note that the direct identification $\rho_{SVI} = \rho_{Heston}$ emerges immediately from the symmetry structure of the discriminant $d(p)$ shown in the previous section. $\square$

Proposition 5.3. *Suppose the fitted surface’s long-end slice is given by the SVI limit (15), yielding the observed parameters $(\omega_1, \omega_2, \rho)$. The underlying structural parameters are recovered as follows: the correlation ρ is identified with the fitted slice’s SVI correlation, and*

$$\kappa\theta = \frac{4}{(1 - \rho^2)\omega_1\omega_2^2}, \quad \xi = \frac{4}{(1 - \rho^2)\omega_1\omega_2}. \quad (17)$$

Furthermore, we now have two redundant readings of κ . One from the at-the-money skew term structure (Theorem 4.4, Lemma 4.5), and one via the product $\kappa\theta$ of (17) divided by the level asymptote θ of Proposition 4.2.

Proof. The exact identification of ρ is a structural consequence of the limiting smile’s asymmetry. Under the Legendre transform of the CGF, the difference between the left- and right-wing asymptotic slopes is strictly governed by the Heston correlation parameter. Because this asymmetric tilt maps directly to the linear ρx term in the SVI parametrization (15), the observed SVI correlation is structurally locked to the Heston ρ .

To invert the remaining scale parameters, we construct the product $\omega_1\omega_2^2$ from equations (16):

$$\omega_1\omega_2^2 = \left(\frac{4\kappa\theta}{\xi^2(1 - \rho^2)} \right) \left(\frac{\xi^2}{(\kappa\theta)^2} \right) = \frac{4}{(1 - \rho^2)\kappa\theta}.$$

Rearranging this isolates the drift product $\kappa\theta$. Substituting this result into the definition of the second SVI parameter, $\omega_2 = \xi/(\kappa\theta)$, yields $\xi = \omega_2\kappa\theta$, which directly produces the final equality. $\square$

With the inversion of the long-maturity slice, the primary objective of our sequential recovery map is achieved. The correlation ρ and volatility-of-volatility ξ are definitively decoupled via the large-maturity asymptotics. The full Heston parameter set $(v_0, \theta, \kappa, \rho, \xi)$ is now uniquely identified through explicit structural features of the variance surface and its SVI manifold.

However, the large-maturity regime contains one final structural constraint. By pushing the spatial smile to its extreme-strike limits (i.e., the wings) we can establish a direct link between the SVI asymptotic slopes and the model-independent moment bounds established by Lee [13]. This final step provides a further consistency check on parameter combinations, adding an item to the over-identified network of Section 7.1.

5.3 Wings

To formalize the aforementioned consistency check, we recall that the asymptotic slopes of the implied variance at extreme strikes are dictated by the maximal finite moments of the

return distribution [13]. In the Heston framework, these critical moments correspond exactly to the boundaries of the CGF domain derived in Proposition 5.1, where the discriminant $d(p)$ vanishes. This is developed in the next proposition.

Proposition 5.4. *The limiting SVI slice (15) exhibits asymptotic annualized-variance wing slopes in the scaled log-moneyness x of*

$$\frac{1}{2}\omega_1\omega_2(1 \pm \rho)$$

as $x \rightarrow \pm\infty$. In the unscaled log-moneyness k the total-variance slopes are the same numbers. These slopes coincide with the theoretical bounds prescribed by Lee's moment formula [13] evaluated at the critical exponents $p_{\pm}$ (the roots of $d(p) = 0$ from Proposition 5.1). Consequently, the SSVI wing slopes, evaluated at large ϑ_T , estimate the same Heston parameter combinations as the long-maturity inversion (17), providing a redundant structural evaluation.

Proof. The asymptotic slopes of the SVI limit (15) can be evaluated directly. Taking the limit as $x \rightarrow \pm\infty$, the square root is dominated by $\omega_2|x|$, yielding the slopes of the annualized implied variance with respect to $|x|$ as

$$\lim_{x \rightarrow \pm\infty} \frac{1}{|x|} \left(\frac{w(xT, T)}{T} \right) = \frac{\omega_1}{2} (\pm \omega_2 \rho + \omega_2) = \frac{1}{2} \omega_1 \omega_2 (1 \pm \rho),$$

and since $w(xT, T)/T$ per unit x equals $w(k, T)$ per unit k , the same numbers are the total-variance wing slopes in k . This is the form in which Lee's bounds and the SSVI wings (6) are stated. This algebraic limit explicitly isolates the combinations required for our recovery map and matches the large- ϑ_T limits of the SSVI parametrization.

To establish that these slopes are consistent with the underlying Heston dynamics, we rely on the model-free moment bounds of Lee [13]. Lee's formula dictates that extreme-strike total-variance slopes, denoted $\beta_{\pm}$, are strictly determined by the critical moments $p_{\pm}$ where the return's moment generating function diverges,

$$\beta_{\pm} = 2 - 4(\sqrt{p_{\pm}^2 \mp p_{\pm}} - p_{\pm}).$$

As demonstrated by Gatheral and Jacquier [9], these critical moments correspond exactly to the boundaries of the Riccati domain ($d(p) = 0$) established in Proposition 5.1. Their derivation algebraically verifies that evaluating Lee's relation at these specific Heston roots produces $\frac{1}{2}\omega_1\omega_2(1 \pm \rho)$. Thus, the SVI wings intrinsically satisfy the theoretical moment bounds, confirming that they provide a redundant read of the parameter groupings in (17). $\square$

Theorem 5.2 imposes a strict structural constraint on any Heston-consistent SSVI parametrization. Matching (5) to (15) at large T requires $\vartheta_T \sim \omega_1 T$ and $\varphi(\vartheta_T) = \omega_2/T \sim \omega_1\omega_2/\vartheta_T$. Because the scale parameter $\omega_2 = \xi/(\kappa\theta)$ is independent of maturity, the SSVI function $\varphi(\vartheta) = \eta\vartheta^{-\gamma}$ must scale as ϑ^{-1} along the term structure, and this Heston limit forces the power-law exponent to $\gamma \rightarrow 1$.

However, equity option markets systematically violate this constraint. Robust surface fits for single-name equities and indices consistently yield power-law exponents in the range of $\gamma \approx 0.4$ to 0.6 [10, 11]. It is tempting to interpret this as a direct measurement of the variance process proper. Specifically, that it is a signature of rough volatility, where skew decays as

$T^{H-1/2}$ with Hurst parameter H near zero [11] since in the power-law coordinates the at-the-money skew decays as $T^{-\gamma}$. The exponent alone does not support that interpretation, however.

A single power law affects *both* ends of the calibration window, and the two ends pull in opposite directions. The Heston skew of Theorem 4.4 is bounded at the short end, with $g(0^+) = \frac{1}{2}$, and decays as $1/T$ at the long end. On the other hand, $T^{-\gamma}$ with γ near one half diverges at the short end and decays more slowly than $1/T$ at the long end. Fitted over a finite maturity window, the exponent is a compromise between the two that is steeper than Heston at short maturities and flatter at long ones. We consider this directly in Section 8 by exhibiting surfaces generated by Heston itself, with no roughness in them and fitted with γ between 0.39 and 0.70 (Table 2).

The tension between the fitted γ and the Heston value $\gamma = 1$ is therefore real, but it is not a function entirely of the exponent. It is the tension between a short-end skew that keeps steepening as $T \rightarrow 0$ and one that levels off at $\frac{1}{2}\rho\xi$. It is resolved at the front of the skew term structure measured slice by slice (Definition 6.1) where the two differ most. Ultimately, the global power law, being a compromise across the entire calibration window, cannot make any distinction, but the per-slice read can, and the over-identification statistics of §7.1 identify this directly.

When evaluated on a simulated surface generated by Heston dynamics, the recovery map inverts the underlying parameters to the precision of the finite-window assumptions (§8). Consequently, when the map is applied to a real market surface, any resulting gaps among its redundant parameter reads are, beyond that precision, not calibration errors. They are quantitative measurements of the model family’s structural misspecification. By leveraging the over-identification network of §7.1, our framework quantifies this misspecification, replacing the boundary behavior of a least-squares optimizer with a reported statistic.

6 The recovery map

With the theoretical boundaries of the Heston surface fully characterized across the small-time, at-the-money, and long-maturity regimes, we now synthesize these isolated asymptotic limits into a unified calibration algorithm. The objective of this section is to formalize the recovery map of the paper as a deterministic, sequential procedure that translates a fitted SSVI surface into the full Heston parameter vector $(v_0, \theta, \kappa, \rho, \xi)$ without reliance on multidimensional search heuristics.

As is evidenced in the preceding sections, this inversion must proceed in a strict sequence due to the demonstrated structural dependencies of the Heston model. We first read the correlation ρ directly from the global SSVI shape and extract the spot variance v_0 from the short-term variance level. With these anchors in place, the term structure of the skew uniquely isolates the mean reversion κ purely through its decay shape, which subsequently unlocks the volatility-of-volatility ξ from the short-end scale. Finally, the long-run variance θ is recovered via a direct linear projection of the variance term structure.

Before proceeding, we must formalize how the correlation parameter is bridged across the calibration window. The short-time regime strictly identifies the combined product $\rho\xi$ (Proposition 3.3), while the Gatheral–Jacquier limit isolates ρ exactly at large maturities (Theorem 5.2). To link these regimes, we rely on the native geometry of the SSVI measurement layer.

Assumption 2. *The SSVI parameterization enforces a constant global correlation ρ_s across the entire volatility surface. We assume this global shape parameter identically represents the Heston correlation ρ and not just at the long-maturity asymptote. This identification acts as a structural gauge condition, propagating the isolated ρ from the large-maturity limit backward across all tenors and uniquely breaking the (ρ, ξ) degeneracy left free by the finite-tenor variance and skew curves.*

This differs in kind from the reads of Section 4. There, the interior tenors sample a single entire function of the model, $\rho\xi g(\kappa T)$, and analyticity carries its boundary values inward. Here, no such function like $g(\cdot)$ exists. At first order in ξ , every finite tenor yields the coupled product $\rho\xi$ but cannot isolate ρ itself, which only emerges from the wings in the asymptotic limit of Theorem 5.2. The surface therefore does not determine the split of the product on any finite window on the interior, and the assumption is therefore necessitated.

The impact of this assumption is a finite- T bias in the recovered $\hat{\rho}$, however. We quantify this bias in §8. The theoretical alternative would be a second-order read of ξ^2 from the at-the-money curvature, at which point the equivalence $\rho_s = \rho$ becomes an over-identification statistic rather than an assumption. We do not pursue this direction in the present work.

We next begin with definitions of specific analytic inputs extracted from an SSVI manifold.

Definition 6.1 (SSVI functionals). *Given a fitted SSVI surface $(\rho_s, \eta, \gamma; \{\vartheta_T\})$ and a tenor window $[T_f, T_m]$, we define, for the SSVI manifold,*

$$\begin{aligned} (\text{level curve}) \quad \mathcal{L}(T) &= w(0, T) & (\text{at-the-money total variance}), \\ (\text{skew curve}) \quad \varsigma_s(T) &= \frac{\rho_s \varphi_T \vartheta_T}{T} & (\text{from (6), } \varphi_T \text{ read slice by slice with } \rho_s \text{ shared}), \\ (\text{long end}) \quad (\omega_1, \omega_2) & & (\text{wing/limit where supported by } [T_f, T_m]). \end{aligned}$$

The skew curve is constructed by pairing the global correlation ρ_s with a tenor-specific φ_T , where each φ_T is independently calibrated to its respective slice via a one-dimensional least-squares fit and not from the global power law $\varphi(\vartheta) = \eta\vartheta^{-\gamma}$. As noted previously, the power law's implied skew curve, $\rho_s \eta \vartheta_T^{1-\gamma}/T$, has the shape of $T^{-\gamma}$ to leading order by construction. This diverges at the short end where the Heston skew is finite (Proposition 3.3), and it replaces the decay shape that isolates κ with a single exponent γ . Section 8 quantifies this discrepancy.

Theorem 6.2 (Quasi-explicit recovery). *Under the asymptotic identifications of Sections 3–5 and Assumptions 1 and 2, the full set of Heston parameters is recovered from the SSVI functionals by:*

$$\begin{aligned} \rho &= \rho_s & (\text{Assumption 2; Thm. 5.2 at the long end}), \\ v_0 &= \lim_{T \rightarrow 0} \mathcal{L}(T)/T \approx \mathcal{L}(T_f)/T_f & (\text{Prop. 3.2}), \\ \kappa &= \arg \min_{\kappa} \min_c \|\varsigma_s - c g(\kappa \cdot)\|_{[T_f, T_m]} & (\text{Thm. 4.4, Lemma 4.5}), \\ \rho\xi &= \varsigma_s(T_{ref}) / g(\kappa T_{ref}) \Rightarrow \xi = (\rho\xi)/\rho_s & (\text{Thm. 4.4}), \\ \theta &= \theta^*(\kappa) = \frac{\langle \mathcal{L} - v_0 B, T - B \rangle}{\langle T - B, T - B \rangle} & (\text{Lemma 4.1, Prop. 4.2;}) \\ & & \text{linear projection at the recovered } \kappa), \end{aligned} \tag{18}$$

where $\langle \cdot, \cdot \rangle$ is any fixed inner product on the calibration window, and $B = B(\kappa, T)$.

The family $\{cg(\kappa \cdot)\}$ is identifiable on the window (Lemma 4.5) and a minimizer of (18) exists on any compact κ -domain.

Proof. Each line corresponds to the associated proposition or theorem, evaluated directly on the SSVI manifold. This substitution is rigorous because ς_s and $\mathcal{L}$ are exact functionals of the fitted surface, which serves as the arbitrage-free measurement layer.

To see (18), the mapping $\kappa \mapsto (g(\kappa T))_{T \in [T_f, T_m]}$ modulo scale is injective by Lemma 4.5, so a skew curve that lies in the family determines κ uniquely, and the existence of the L^2 minimizer on a compact κ -domain is standard. The skew curve of Definition 6.1 is read slice by slice and does not in general lie in the family, though. As a result, (18) is a projection, and injectivity of the family does not by itself make its minimizer unique.

The conditioning of this projection is strictly governed by the curvature of its L^2 objective profile. A locally flat objective does not indicate a computational failure of the map, but rather generates a direct specification statistic. By reporting this curvature alongside the parameter gaps within the over-identification network of §7.1, the profile explicitly quantifies when the finite calibration window lacks sufficient maturity depth to resolve the skew decay. This measurement captures the exact same structural degeneracy that produces the ill-conditioned (κ, θ) calibration ridge on the level curve (Remark 4.3).

The factorization of $\rho\xi$ relies on the identification $\rho = \rho_s$ of Assumption 2, which bridges the two boundary regimes. Namely, Proposition 3.3 identifies the product at the short end, while Theorem 5.2 isolates ρ itself at the long end. Because the SSVI parameterization carries a single global ρ_s across the entire surface, this decoupled recovery is exact under the assumption. $\square$

We have arrived now at the main claim of the work. Each step of the Heston parameter identification evaluates an explicit functional or solves a scalar problem, and the recovery map involves no multidimensional optimization.

While Theorem 6.2 secures parameter recovery for fully featured volatility surfaces, we next consider calibration on degenerate market regimes. Standard numerical optimizers are notoriously brittle when presented with flat variance term structures or absent volatility smiles where the lack of geometric features flattens the objective function. Frequently, this causes unconstrained solvers to push parameters to preset boundaries which destabilizes subsequent Greeks.

Next we show that the deterministic map established here is structurally continuous into its own boundary. When confronted with featureless market data, the sequential inversion does not fail computationally. Instead, it natively identifies the degeneracy and exactly reduces the Heston dynamics to their correct limiting sub-models.

Proposition 6.3. *The recovery map inherently accommodates degenerate surface limits without estimator failure. We have*

- (i) **Flat level curve:** *If the at-the-money variance is strictly linear in time such that $\mathcal{L}(T) = v_0 T$ across the calibration window, the long-run variance recovers as $\theta^*(\kappa) = v_0$ identically for any choice of κ . The mean reversion parameter κ becomes structurally inert.*
- (ii) **Flat smile:** *If the volatility smile exhibits no spatial variation across strikes such that $\varphi(\vartheta_T) \equiv 0$ across the calibration window, then the recovery map returns $\xi = 0$. The model*

collapses to a deterministic-volatility Black-Scholes framework, rendering the correlation ρ irrelevant.

Proof. We evaluate the explicit limits of the sequential map in Theorem 6.2.

Part (i): Suppose the term structure exhibits a flat local variance, meaning the total integrated variance grows strictly linearly: $\mathcal{L}(T) = v_0 T$. We evaluate the linear projection for θ by substituting this directly into the numerator:

$$\begin{aligned}\theta^*(\kappa) &= \frac{\langle v_0 T - v_0 B(\kappa, T), T - B(\kappa, T) \rangle}{\langle T - B(\kappa, T), T - B(\kappa, T) \rangle} \\ &= v_0 \frac{\langle T - B(\kappa, T), T - B(\kappa, T) \rangle}{\langle T - B(\kappa, T), T - B(\kappa, T) \rangle} \\ &= v_0.\end{aligned}$$

This algebraic cancellation holds exactly for *every* value of $\kappa > 0$. Because the variance level is constant, the mean reversion speed has no physical manifestation in the data. The parameter κ is entirely unidentified but physically inert. Fixing κ at *any* conventional positive value alters no calculated quantity and will produce the correct flat-variance dynamics.

Part (ii): Suppose the volatility surface is flat across strikes, so that the SSVI spatial scale satisfies $\varphi(\vartheta_T) \equiv 0$ across the calibration window. By (5), the total implied variance collapses to $w_{\text{SSVI}}(k, \vartheta_T) \equiv \vartheta_T$ for all k . This eliminates all strike dependence, skew, and smile curvature.

From Definition 6.1, the skew curve vanishes identically ($\varsigma_s(T) \equiv 0$), which anchors $\rho\xi = 0$. Furthermore, the absence of asymptotic wing divergence in (6) forces the SVI scale parameter $\omega_2 \rightarrow 0$. Via the long-maturity inversion (17), $\xi = \omega_2 \kappa \theta = 0$, uniquely resolving the indeterminacy and forcing the volatility-of-volatility to vanish.

Substituting $\xi = 0$ into the Heston system (1)–(2), the diffusion term of the variance equation drops out, leaving v_t to evolve along the deterministic trajectory $dv_t = \kappa(\theta - v_t) dt$. Because the return process dX_t no longer couples with a stochastic variance driver dZ_t , the Brownian motion dZ_t ceases to exist in the dynamics, rendering the leverage correlation $\rho = d\langle W, Z \rangle_t / dt$ meaningless. The recovery map thus safely collapses the Heston geometry into a Black-Scholes model without numerical singularity. $\square$

7 Practical Implementation and Model Specification

The quasi-explicit recovery map both guarantees computational stability of a Heston parameter determination and highlights the inherent limitations of the underlying stochastic geometry. In practical implementation, this framework introduces three critical advantages for quantitative specification.

7.1 The over-identification network

Because the Heston parameter vector is fully determined before all available surface information is exhausted, the map naturally recovers several parameters more than once. For instance, the mean reversion κ is identified from the skew decay (18), implied by the variance level blend (8), and extractable from the long-maturity wings (Proposition 5.3). Similarly, ξ is seen through the skew scale, the small-time curvature (Proposition 3.4), and the long-end inversion.

In a strictly theoretical Heston environment, all such redundant assessments will agree perfectly. When applied to empirical market data, however, their pairwise gaps become direct *specification statistics* that are continuous, interpretable functionals measuring where a single-factor square-root variance process fails to reconcile the surface’s competing regimes (as anticipated in our discussion of the power-law exponent closing §5.3). The profile of the κ criterion (18) belongs to the same network. Its flatness across the domain is the residue of the level ridge and is reported beside the gaps. This network of measurements completely replaces the opaque detection-and-fallback logic required by ridge-afflicted least-squares solvers, for example.

7.2 Finite-tenor evaluation

The theoretical results underpinning this map are exact in their respective asymptotics ($T \rightarrow 0$, first-order in ξ , and $T \rightarrow \infty$). In a production environment, however, the map may evaluate these functionals at finite tenors bounded within a defined calibration window.

The induced approximation bias is simply the difference between the model’s finite- T smile and its asymptotic form. Because our methodology immediately yields the Heston parameters without iterative search, this bias is readily computable by pricing the surface under the recovered parameters and differencing the result against the fitted SSVI layer. Consequently, the approximation cost of the entire construction may be directly reported and observed.

7.3 Calibration to fixed-delta instruments and the skew adjustment

Finally, we address the practical reality of surface construction and hedging. It is often desirable to anchor the variance level data $\mathcal{L}$ not at the money, but along a highly liquid fixed-delta slice (for instance, the out-of-the-money 15Δ or 25Δ put). This ensures that the calibrated Heston dynamics strictly respect the primary traded hedging instrument.

The sequential recovery map naturally accommodates this fixed-delta calibration, provided the structural impact of the skew is explicitly decoupled from the variance level. When the level curve is evaluated along a slice with non-zero log-moneyness $k(T)$, the total implied variance expands via Theorem 4.4 as

$$w_{\text{slice}}(T) = \underbrace{v_0 B(\kappa, T) + \theta(T - B(\kappa, T))}_{\text{baseline variance level}} + \underbrace{T \varsigma(T) k(T)}_{\text{first-order skew offset}},$$

where the skew is given by $\varsigma(T) = \rho\xi g(\kappa T)$.

To extract the correct fixed-delta dynamics, the recovery sequence is executed with a single, explicit adjustment. The algorithm proceeds by first identifying $(\kappa, \rho\xi)$ directly from the skew curve. Using these parameters, it computes the first-order skew offset, $T \varsigma(T) k(T)$, and subtracts it from the observed traded variance, $w_{\text{slice}}(T)$. The v_0 short-time evaluation and the θ linear projection from Theorem 6.2 are then applied exclusively to this *skew-adjusted* level curve.

If one instead naively matches the raw out-of-the-money traded slice directly to the baseline variance equation, the calibration severely double-counts the skew premium. It absorbs the skew once statically into the baseline parameters (v_0, θ) , and then generates it a second time dynamically through the correlation mechanics of (ρ, ξ) . *The resulting model will systematically overprice the target traded slice.* This structural error is frequently masked in standard

optimizers because they tend to artificially understate ξ , but it becomes glaringly apparent the moment our sequential map restores ξ to its structurally identified magnitude.

Because the skew-adjusted curve theoretically coincides with the ATM level curve at first order, all mathematical statements regarding Theorem 6.2 and the degenerate limits of Proposition 6.3 apply unchanged. In the deterministic Black–Scholes reduction, the skew offset naturally vanishes alongside $\rho\xi$, ensuring that data primacy remains strictly with the fixed-delta instrument while maintaining theoretical consistency.

8 Numerical Recovery

To empirically validate the recovery map, we evaluate its performance on a panel of simulated surfaces with known parameters. Following the protocol of the quasi-explicit SVI calibration [16], we generate the theoretical surface, apply the intermediate SSVI fit, and execute the sequential recovery map. The recovered parameters are then benchmarked under both exact and noisy conditions. Finally, we systematically sweep the two primary sources of theoretical bias identified in our derivation: the relative gap $(v_0 - \theta)/\theta$ of Assumption 1 and the truncation of the tenor window (Remark 4.3).

8.1 Simulation and Recovery Protocol

Data Generation and the Measurement Layer. Heston prices are generated via the Lewis integral representation of the characteristic function and inverted into implied variances. The simulation grid comprises eight tenors ranging from one month to two years ($T_{\max} = 2$), with eight strikes per tenor defined by the scaled log-moneyness $k = z \sigma_{\text{ATM}} \sqrt{T}$ for $z \in \{-2, -1.5, \dots, 1.5\}$. To construct the measurement layer, a global SSVI surface (ρ_s, η, γ) is fitted to each generated grid using the at-the-money total variance ϑ_T per slice as data.

Execution of the Map and Finite-Tenor Corrections. The sequential recovery map of Theorem 6.2 is then executed, incorporating two native structural corrections. First, the mean reversion incorporates the second-order de-biasing $\kappa = \kappa_{\text{eff}} + \rho\xi/2$. Second, the map applies a fixed-point correction to the level curve to account for finite-tenor biases. Specifically, the model’s theoretical gap $\mathcal{W}(T) - w(0, T)$ (the $O(\xi^2)$ perturbation of Proposition 4.2) is computed at the provisionally recovered parameters and added to the level curve prior to the θ projection. Inside this same fixed point, the short-end read $v_0 = \mathcal{L}(T_f)/T_f$ is back-solved through the $\mathcal{W}$ law at the recovered (κ, θ) to remove its $O(T_f)$ contamination. This joint fixed-point iteration reliably converges in three steps.

Test Panel. The map is evaluated across six distinct parameter vectors (Table 2) chosen to span a spectrum of market regimes: a standard equity index, a characteristic single-name equity, slow and fast mean reversion profiles, a flat forward variance curve where Assumption 1 is structurally exact, and a diffusion strictly violating the Feller condition ($\xi^2 > 2\kappa\theta$). As an illustrative baseline, Figure 2 displays the single-name surface across maturities, overlaying the global SSVI fit and the Heston smile repriced at the recovered parameters.




Figure 1: Term structure dynamics for the characteristic single-name surface (Vector B in Table 2). **Left panel:** The at-the-money variance skew across maturities. Markers denote the per-slice empirical skew. The solid curve represents the sequentially fitted theoretical shape $\hat{\rho}\hat{\xi}g(\hat{\kappa}_{\text{eff}}T)$, which strictly isolates the mean reversion κ through its maturity decay profile and the scale product $\rho\xi$ at the short end. The dashed curve traces the global power-law skew $(\rho_s\eta\vartheta_T^{1-\gamma}/T)$, explicitly illustrating its divergence as $T \rightarrow 0$. **Right panel:** The annualized at-the-money implied variance ϑ_T/T (markers) contrasted with the integrated variance expectation $\mathcal{W}(T)/T$ evaluated at the exact baseline (dotted) and recovered parameters (solid). As established in Proposition 4.2, the implied variance exhibits a structural downward perturbation of $O(\xi^2)$ relative to $\mathcal{W}(T)/T$ (here approximately 0.01 in annualized variance for $\xi = 1.22$). This finite-tenor gap is rigorously eliminated via the fixed-point repricing correction.

8.2 The measurement layer

Before evaluating the full recovery map, we isolate the structural input to the mean reversion parameter. Table 1 applies the κ projection to three distinct formulations of the skew curve:

- The global power law's implied skew $\rho_s\eta\vartheta_T^{1-\gamma}/T$.
- The surface's true at-the-money skew calculated via central difference.
- The slice-by-slice SSVI evaluation formalized in Definition 6.1.

When extracted from the global power law, the recovered $\hat{\kappa}$ is severely upward-biased, exceeding the true value by a factor of three to five across all test surfaces, including Vector E where Assumption 1 holds exactly. The corresponding $\hat{\xi}$ is roughly doubled. This failure is a direct consequence of the term-structure geometry. As noted previously, the global power law forces the skew term structure to decay as $T^{-\gamma}$. With fitted exponents in the range $\gamma \in [0.39, 0.70]$, this functional form artificially diverges at the short end where the true Heston skew is rigidly bounded (Proposition 3.3). To compensate, the L^2 projection against $g(\kappa T)$ forces a spuriously large κ . Consequently, fitting against the global power law acts as a projection of γ and not a statement on mean reversion.

In contrast, when κ is extracted from the true surface skew or the per-slice φ_T , $\hat{\kappa}$ recovers to within 10–25% for moderate ξ , while $\hat{\xi}$ is identified within 5%. Thus, while the global

vector	$\hat{\kappa}$				$\hat{\xi}$			
	true	power law	surface	per slice	true	power law	surface	per slice
A	2.00	8.52	2.23	2.17	0.30	0.64	0.31	0.30
B	3.87	19.87	5.77	5.02	1.22	3.26	1.26	1.29
C	0.75	4.88	1.62	1.37	0.50	0.98	0.53	0.54
D	6.00	29.32	7.27	6.60	0.80	2.52	0.87	0.84
E	3.00	15.27	3.79	3.47	0.60	1.61	0.63	0.63
F	1.50	15.65	7.36	4.81	1.50	3.52	1.52	1.66

Table 1: The κ and ξ reads from a single mapping evaluated on three distinct skew curves (noise-free, without level correction). “Surface” denotes the central-difference ATM skew utilizing the true ρ ; “per slice” denotes φ_T fitted independently slice by slice using the global ρ_s .

power law remains the measurement layer for isolating the global correlation ρ_s and for the level curve ϑ_T , it is structurally invalid as the measurement layer for the skew term structure. This dichotomy explicitly dictates the per-slice evaluation mandated in Definition 6.1. Figure 1 visually contrasts these term structures on a characteristic single-name surface.

8.3 Exact recovery

Table 2 reports the map executed under the per-slice read with the fixed-point level correction on noise-free data.

vector	κ	$\hat{\kappa}$	θ	$\hat{\theta}$	ξ	$\hat{\xi}$	ρ	$\hat{\rho}$	v_0	$\hat{v}_0$	γ_{fit}
A	2.00	2.17	0.040	0.040	0.30	0.30	−0.70	−0.69	0.040	0.040	0.50
B	3.87	5.02	0.097	0.093	1.22	1.29	−0.32	−0.29	0.082	0.083	0.67
C	0.75	1.37	0.060	0.048	0.50	0.54	−0.60	−0.52	0.040	0.041	0.39
D	6.00	6.60	0.050	0.050	0.80	0.84	−0.50	−0.48	0.090	0.091	0.70
E	3.00	3.47	0.060	0.059	0.60	0.63	−0.40	−0.37	0.060	0.060	0.61
F	1.50	4.81	0.050	0.032	1.50	1.66	−0.20	−0.15	0.050	0.054	0.54

Table 2: Exact recovery utilizing the per-slice skew read and level correction for $T_{\max} = 2$. A: equity index. B: characteristic single-name equity. C: slow reversion, $\kappa T_{\max} = 1.5$. D: fast reversion, inverted level curve. E: flat forward variance. F: Feller violated, $\xi^2 = 2.25\kappa\theta$.

For Vectors A, D, and E, the recovery is highly accurate: $\hat{\kappa}$ is isolated to within 9–16%, $\hat{\theta}$ to within 2%, $\hat{\xi}$ to within 6%, and $\hat{v}_0$ to within 2%. Every observed deviation from the exact parameters is structurally predicted by the asymptotic theory. For instance, the recovered correlation $\hat{\rho}$ consistently attenuates toward zero by 2–14% of $|\rho|$ (and by 26% on Vector F, where the absolute correlation is inherently small). This attenuation explicitly quantifies the finite- T bias incurred by Assumption 2. A single global ρ_s is fitted across tenors that have not yet converged to the Gatheral–Jacquier large-maturity limit.

The precision of the spot variance relies heavily on the fixed-point correction at the front of the curve. Without back-solving the initial short-end read, $\hat{v}_0$ would exhibit a downward bias of 2–14%, peaking in Vector D where κT_f and $|v_0 - \theta|$ are largest. This bias corresponds exactly to the $O(T_f)$ correction identified in Proposition 3.2. Because this is generated entirely

Single-name surface: $\kappa = 3.87$, $\theta = 0.097$, $\xi = 1.22$, $\rho = -0.32$, $v_0 = 0.082$; recovered $\hat{\kappa} = 5.20$, $\hat{\theta} = 0.092$, $\hat{\xi} = 1.31$, $\hat{\rho} = -0.29$, $\hat{v}_0 = 0.082$




Figure 2: Implied volatility versus log-moneyness for the characteristic single-name surface (Vector B in Table 2) across six maturities. Markers denote the simulated surface perturbed by 20 basis points of Gaussian noise. Dashed lines represent the intermediate global SSVI fit, while solid lines depict the Heston model repriced strictly at the recovered parameters. The global power-law fit structurally overstates the one-month put wing, manifesting the $T^{-\gamma}$ divergence inherent in its skew term structure. By contrast, the recovered Heston smile aligns with the simulated surface to within 50 basis points of implied volatility for tenors $T \geq 0.25$, and within 100 basis points at the one-month boundary.

by the deterministic drift of the $\mathcal{W}$ expectation law, a single back-solve against the recovered drift parameters provides the needed correction.

The remaining vectors highlight the map's behavior at the boundaries of its theoretical validity. On Vector C, the mean reversion is slow relative to the calibration window ($\kappa T_{\max} = 1.5$). Consequently, the theoretical shape $g(\kappa T)$ is nearly flat across the available tenors, producing a similarly flat projection profile for κ . This forces an 80% upward bias in $\hat{\kappa}$, which subsequently drags $\hat{\theta}$ downward by 21%. This is the exact structural degeneracy anticipated in Remark 4.3, now manifesting in the skew read because the truncated window fails to capture the decay of either curve. Finally, Vector F pushes the first-order skew theorem far outside its perturbative regime ($\xi = 1.5$). While this inflates $\hat{\kappa}$ by a factor of three, the volatility-of-volatility $\hat{\xi}$ is still recovered to within 11%, because its scale is anchored at the short end where the first-order expansion remains tight. The 8% upward bias in $\hat{v}_0$ is a cascading artifact, as its back-solve is executed using the distorted κ .

8.4 Noise and empirical stability

To assess the robustness of the recovery map, we perturb each simulated implied volatility with independent Gaussian noise of 20 basis points. The map is then executed across 100 replications per surface, with the resulting sample distributions summarized in Table 3.

vector	$\hat{\kappa}$	s.d.	$\hat{\theta}$	s.d.	$\hat{\xi}$	s.d.	$\hat{\rho}$	s.d.
A	2.16	0.17	0.0395	0.0006	0.298	0.021	-0.691	0.053
B	5.07	0.33	0.0925	0.0009	1.294	0.049	-0.290	0.010
C	1.36	0.10	0.0478	0.0009	0.544	0.015	-0.514	0.010
D	6.60	0.41	0.0499	0.0007	0.845	0.048	-0.478	0.031
E	3.45	0.28	0.0590	0.0007	0.632	0.031	-0.370	0.019
F	4.82	0.23	0.0318	0.0005	1.662	0.039	-0.148	0.005

Table 3: Sample mean and standard deviation of the recovered parameters over 100 replications, with 20 basis points of independent Gaussian noise applied to the implied volatilities. The map utilizes the per-slice read and fixed-point level correction with $T_{\max} = 2$.

The resulting standard deviations are small in absolute terms, with the long-run variance θ stable to four decimal places, ξ and ρ to a few hundredths, and κ to 5–8%. Where Table 2 showed a material bias (κ on Vectors B, C and F; θ on C and F), that bias exceeds the noise dispersion by a factor of three to twenty-five. Where it showed little bias (θ , ξ and ρ on Vectors A, D and E), the dispersion is the whole of the error, and for κ on those vectors bias and noise are of comparable size.

Where a bias is present it therefore dominates the noise, and the biases themselves do not move under noise. This is the property that matters for a daily calibration: the recovered parameters are stable from day to day, and where they are wrong they are wrong by a known amount in a known direction.

8.5 The two named biases

Two biases are identified in the derivation of the recovery map. First, the level term structure entering the skew through Assumption 1, and second, the truncation of the tenor window in Remark 4.3. Table 4 systematically sweeps the variance level term structure at fixed $(\kappa, \theta, \xi, \rho) = (2, 0.04, 0.3, -0.7)$, and subsequently sweeps the calibration window length at fixed $(\kappa, \theta, \xi, \rho, v_0) = (1, 0.04, 0.3, -0.7, 0.06)$.

The bias induced by Assumption 1 is strictly monotone in the relative gap $(v_0 - \theta)/\theta$, moving $\hat{\kappa}$ from +13% at a gap of -0.5 to -7% at a gap of 2, for a swing of 20%. A single fixed-point evaluation of the generalized shape (12) at the provisionally recovered (v_0, θ) reduces that swing to under 4%. This is the scalar correction outlined following Theorem 4.4. The residual offset of approximately 8% is the same one Vector A exhibits at a flat curve (Table 2), and is therefore not attributable to the level term structure.

The window-length sweep highlights the remedy proposed in Remark 4.3. By extracting κ from the skew shape and back-solving the front-end v_0 , the recovered $\hat{\kappa}$ remains within 11% and $\hat{\theta}$ within 2% for all windows ranging from $\kappa T_{\max} = 0.5$ to 5.0. Attempting to identify both (κ, θ) simultaneously from a truncated level curve is an ill-conditioned inverse problem. However, once κ is anchored by the skew term structure, the θ projection becomes well-conditioned, even on windows that terminate long before the variance decay is complete. Mathematically, this

$(v_0 - \theta)/\theta$	$\hat{\kappa}$	$\hat{\kappa}$, one step of (12)	$\hat{\theta}$	$\hat{\xi}$
-0.50	2.26	2.20	0.039	0.289
-0.25	2.23	2.20	0.039	0.297
0	2.17	2.19	0.040	0.300
0.25	2.12	2.18	0.040	0.302
0.50	2.07	2.17	0.040	0.302
1.00	1.98	2.15	0.040	0.302
2.00	1.86	2.13	0.039	0.300
$T_{\max} (\kappa T_{\max})$	tenors	$\hat{\kappa}$	$\hat{\theta}$	$\hat{\xi}$
0.5	4	1.07	0.039	0.280
1.0	6	1.09	0.040	0.296
2.0	8	1.10	0.040	0.306
3.0	9	1.11	0.040	0.308
5.0	11	1.11	0.040	0.309

Table 4: Upper panel: The recovered $\hat{\kappa}$ as the initial variance term structure departs from flatness (Assumption 1), shown both raw and after a single fixed-point evaluation of the generalized first-order shape (12). Lower panel: Parameter recovery across varying calibration window lengths $T_{\max}$ evaluated at $\kappa = 1$.

stability arises because the projection now merely executes a linear regression of the level curve against the established basis function $T - B(\kappa, T)$. This calibration ridge only resurfaces if the skew curve itself lacks sufficient data to resolve the curvature; e.g., the structural degeneracy observed in Vector C, where the mean reversion is so slow that $g(\kappa T)$ remains essentially flat across the available tenors.

8.6 Synthesis

The deterministic map successfully recovers the underlying parameters of a Heston surface to the precision the assumptions allow, and each residual error is attributable to an explicit structural assumption. These residual errors are shown to be in line with both theoretically predicted magnitude and direction.

- The finite- T correlation gauge dictates the bias in $\hat{\rho}$.
- The first-order skew truncation and the level term structure dictate $\hat{\kappa}$.
- The window truncation induces a joint degeneracy in $(\hat{\kappa}, \hat{\theta})$ when the true reversion is slow relative to the observable tenors.

Furthermore, because the map natively delivers the parameters necessary to compute the theoretical offsets, the $O(T_f)$ impact in the short-end read and the $O(\xi^2)$ perturbation in the level curve are not merely tolerated as bounds. They are removed by the fixed-point repricing iteration.

The numerical analysis also establishes that the global power law's skew term structure is inadmissible as an input to the κ projection. A calibration that reads a mean reversion parameter from it recovers the exponent γ in place of κ , and inflates ξ along with it. Because the Feller condition compares ξ^2 to $2\kappa\theta$, this paired distortion produces spurious Feller violations.

Any census of violations produced by such a calibration should be reread with that in mind. Notably, as §5 proves and this section demonstrates, the remedy is not a better power law but a different input. That input is the slice-by-slice evaluation of Definition 6.1

9 Conclusion

The recovery map formalizes a strict division of labor. SSVI serves as the estimator—low-dimensional, quasi-explicit, and stable—while Heston is treated as a set of *coordinates* on the SSVI fit, reached through asymptotic correspondences obtained from theorems cut across short and long maturity asymptotics and the term structure holistically.

Three consequences of the recovery map directly impact practice. First, ridge pathologies that otherwise plague least squares fits for calibration become structurally impossible. Second, degenerate surfaces cleanly resolve into exact model reductions (Proposition 6.3) rather than triggering estimator failures that require algorithmic guards. And, third, the model’s genuine finite- T misspecification (such as the power-law and rough-volatility tension discussed in Section 5) is concentrated into a small set of continuous over-identification statistics which can be explicitly monitored. This replaces the standard paradigm where specification error is scattered into optimizer boundary behavior.

Finally, the numerical recovery of §8 adds two statements beyond the proofs elsewhere in the paper. First, every residual error in the map is attributable to a named structural assumption with the magnitude and direction the theory predicts. For instance, the measured attenuation of $\hat{\rho}$ is not an estimator failure, but the exact, anticipated cost of enforcing the finite-tenor gauge condition ($\rho = \rho_s$) rather than formally evaluating the second-order at-the-money curvature. Second, the skew term structure must be evaluated strictly slice by slice. Extracting the mean reversion from a global power law inadvertently recovers the spatial exponent γ in place of κ , simultaneously inflating ξ . Consequently, a calibration built upon a global power law mechanically manufactures Feller violations that the underlying surface does not contain.

References

- [1] L. Andersen, *Simple and efficient simulation of the Heston stochastic volatility model*, Journal of Computational Finance 11 (2008), 1–42.
- [2] H. Berestycki, J. Busca, I. Florent, *Asymptotics and calibration of local volatility models*, Quantitative Finance 2 (2002), 61–69.
- [3] L. Bergomi, J. Guyon, *Stochastic volatility’s orderly smiles*, Risk (May 2012), 60–66.
- [4] D. Duffie, J. Pan, K. Singleton, *Transform analysis and asset pricing for affine jump-diffusions*, Econometrica 68 (2000), 1343–1376.
- [5] M. Forde, A. Jacquier, *Small-time asymptotics for implied volatility under the Heston model*, Int. J. Theor. Appl. Finance 12 (2009), 861–876.
- [6] M. Forde, A. Jacquier, *The large-maturity smile for the Heston model*, Finance and Stochastics 15 (2011), 755–780.
- [7] M. Forde, A. Jacquier, R. Lee, *The small-time smile and term structure of implied volatility under the Heston model*, SIAM J. Financial Math. 3 (2012), 690–708.
- [8] J. Gatheral, *The Volatility Surface: A Practitioner’s Guide*, Wiley, 2006.
- [9] J. Gatheral, A. Jacquier, *Convergence of Heston to SVI*, Quantitative Finance 11 (2011), 1129–1132.
- [10] J. Gatheral, A. Jacquier, *Arbitrage-free SVI volatility surfaces*, Quantitative Finance 14 (2014), 59–71.
- [11] J. Gatheral, T. Jaisson, M. Rosenbaum, *Volatility is rough*, Quantitative Finance 18 (2018), 933–949.
- [12] S. L. Heston, *A closed-form solution for options with stochastic volatility with applications to bond and currency options*, Review of Financial Studies 6 (1993), 327–343.
- [13] R. Lee, *The moment formula for implied volatility at extreme strikes*, Mathematical Finance 14 (2004), 469–480.
- [14] A. L. Lewis, *Option Valuation under Stochastic Volatility*, Finance Press, 2000.
- [15] R. Lord, R. Koekkoek, D. Van Dijk, *A comparison of biased simulation schemes for stochastic volatility models*, Quantitative Finance 10 (2010), 177–194.
- [16] Zeliade Systems, *Quasi-explicit calibration of Gatheral’s SVI model*, Zeliade White Paper (2009).
- [17] S. Mikhailov, U. Nögel, *Heston’s stochastic volatility model: Implementation, calibration and some extensions*, Wilmott Magazine (July/August 2003), 74–79.