

The Two Volatilities

A framework for separating economic volatility from valuation uncertainty

Author: Julius H. Herfel

Date: 6 May 2026

Executive Summary

Financial risk systems often mistake a measurement problem for a market problem. They can conflate two distinct uncertainties: **aleatory**, or aleatoric, uncertainty - the randomness of the world: price movements, earnings surprises, and rate shifts - and **epistemic uncertainty**, which concerns whether our estimate of value is correct, and which arises from model assumptions, stale data, and infrequent transactions. While finance has developed a comprehensive toolkit for managing the former, approaches for managing the latter explicitly are sparse.

The conflation of both concepts is most damaging in illiquid private markets. Model-derived, infrequently tested valuations suppress apparent volatility and can make private assets look safer than they are. Internal economic capital or valuation reserves may be too low where this dynamic goes unrecognised; hedge ratios calibrated on reported marks are liable to track measurement noise; and unmeasured valuation uncertainty accumulates until a liquidity event forces a discontinuous correction.

Epistemic uncertainty requires distinct governance: use valuation ranges instead of single estimates, apply uncertainty overlays or reserves, treat as internal economic capital, and stress-test hedge ratios against desmoothed returns. The core risk framework question is: what uncertainty is being measured?

I. Introduction

Much market practice operationalises volatility as the standard deviation of historical or modelled returns. The definition is so embedded in institutional infrastructure that it has stopped being recognised as a definition at all. It functions as a default convention rather than a deliberate choice, and most of the time the distinction does not matter. When it does matter, it matters substantially.

Financial systems spend enormous energy measuring how prices move. They spend rather less asking how much confidence should be placed in the prices they hold. The tools developed to answer the first question are not fit for purpose for answering the second, and the conflation shows up in capital buffers calibrated against the wrong quantity, in hedge ratios that respond to measurement noise rather than economic exposure, and in private-market portfolios that look stable until the moment they do not.

The title of this note uses “two volatilities” as shorthand. The deeper point is that one of them is not volatility in the conventional market-risk sense at all. One is uncertainty about economic movement; the other is uncertainty about the quality of the valuation through which that movement is observed.

This note argues that financial systems can conflate two fundamentally different kinds of uncertainty: aleatory and epistemic. Separating them is a precondition for a sound valuation infrastructure.

This note does not claim that all private-market valuation stability is artificial, nor that private assets are necessarily mispriced. The narrower claim is that reported volatility from infrequently marked assets cannot, on its own, distinguish economic stability from valuation opacity.

II. The Volatility We Know: Economic Uncertainty

The volatility that finance was built to manage is the randomness of the world. Interest rates rise and fall. Corporate earnings surprise in both directions. Oil prices respond to geopolitical shocks. Currency pairs reprice as inflation differentials widen. Equity indices move with the business cycle. This is the randomness that financial markets were designed to price and redistribute.

Modern finance is comparatively well-equipped to deal with this kind of uncertainty. From Black-Scholes-style option pricing to dynamic delta hedging, and the Harrison-Kreps and Harrison-Pliska martingale foundations of modern asset pricing, financial engineering has produced a sophisticated toolkit for isolating, measuring, and redistributing economic randomness. Quantitative finance over the past half-century has been, in effect, machinery for pricing and redistributing economic risk, and the global system today processes immense volumes of it at low cost and with high speed.

Much of the mainstream toolkit, however, is operationalised through models that often assume estimable distributions, stable parameters, and tractable tail behaviour. One important channel in

the 2007–08 structured credit failure was the fragility of Gaussian-copula correlation assumptions, alongside underwriting failures, incentive misalignments, leverage, rating-agency errors, liquidity evaporation, and housing-market assumptions. The losses were exposed to material underestimation through, among other factors, the dominant modelling architecture and by the governance systems around it.

Economic uncertainty is therefore not the only kind of uncertainty in the financial system. The cases where the toolkit struggles tend to be cases where a different kind of uncertainty dominates.

III. The Factor We Ignore: Epistemic Uncertainty

The second kind of uncertainty is harder to see, precisely because it does not announce itself as a price movement, show up in a time series, or trigger a volatility spike. It arises from incomplete knowledge of where the world currently is, rather than from movement in the world itself.

Consider an illiquid private asset - a stake in a mid-market company, or an infrastructure project with a thirty-year revenue profile. The economic value of that asset changes over time as revenues fluctuate, interest rates move, and the competitive landscape shifts. That is economic uncertainty, and it is real. Sitting on top of it is a second layer: uncertainty about whether the valuation process used to estimate the asset's current worth is reliable. The discount rate is a choice; the revenue forecast embeds assumptions; the comparables that calibrate the multiple were selected by someone against criteria that may no longer hold; and the data underlying the estimate is itself of a particular vintage, which is rarely current.

This is epistemic uncertainty, or valuation uncertainty: uncertainty about the reliability of our estimate of value, distinct from uncertainty about the future movement of value itself.

The substance of the distinction is already encoded in fragments across financial regulation and practice, though each regime is specific to its own context and the parallels are analogical rather than equivalent. On 17 April 2026, the Federal Reserve, the Office of the Comptroller of the Currency and the Federal Deposit Insurance Corporation issued SR 26-2, *Revised Guidance on Model Risk Management*, which superseded and replaced SR 11-7. One foundational tenet of SR 11-7 remains intact in the revised framework: model risk is not merely a technical modelling issue, but a governance and control issue. The EU's prudent valuation regime separates Market Price Uncertainty and Model Risk as distinct Additional Valuation Adjustments deducted from Common Equity Tier 1 capital; that regime is specific to its own prudential context and is referenced here by analogy.

A related but narrower logic appears in the trading-book treatment of risk-factor modellability - the Non-Modellable Risk Factor regime under the FRTB Internal Models Approach - which turns on observable real-price data for individual risk factors and is not itself a private-market valuation framework. In the accounting domain, IFRS 13 and ASC 820 are related but not identical fair-value frameworks: IFRS 13 requires Level 3 sensitivity disclosures, and ASC 820 sets out the US GAAP fair-value hierarchy, with each regime imposing its own disclosure obligations. Underneath all of this sits a substantial academic literature on Knightian uncertainty, robust control, and model uncertainty.

The substance exists. What is missing is consolidation. The aleatory-epistemic distinction is encoded in the regulatory frameworks and academic literatures referenced above, but in fragments. There is no single, unifying architecture and the distinction is not applied consistently across the asset classes - especially where epistemic uncertainty dominates. As a result, the governance and internal capital or reserve implications have not been spelled out for private-market practitioners. These are the gaps this and future research notes address.

None of the underlying components of this argument is new. Getmansky, Lo and Makarov documented return smoothing in 2004; supervisory model-risk governance has existed since SR 11-7 in 2011; the prudent valuation regime has separated market-price uncertainty from model risk since 2016; and the philosophical distinction between aleatory and epistemic uncertainty predates modern finance entirely. The contribution of this note is primarily consolidation. The argument is that these fragments describe a single underlying problem. The problem is especially salient in private markets, where prudential, accounting, conduct and valuation frameworks do not form a single integrated architecture for valuation uncertainty. Practitioners would therefore benefit from treating the distinction as an explicit feature of valuation infrastructure, rather than an implicit one scattered across regimes that were never designed to speak to one another.

Analogy: The Blurred Photograph

Imagine you are observing an object through a camera. The object is moving - slowly, continuously, in ways that are somewhat random but broadly predictable. This movement is the economic volatility of the asset.

The camera, however, is also out of focus. The image on your screen is blurred. The object appears larger or smaller than it is; its edges are uncertain; details that would help you understand its current state are obscured. This is the epistemic uncertainty of the asset. It does not arise from the object moving, but the degraded quality of the channel through which you are observing.

Confuse the two and the consequences follow quickly. If the blur leads you to read more violent motion than is actually present, you will hedge against movements that are not occurring, design responses calibrated to a more chaotic world than the one you inhabit, and consume resources managing a risk that does not exist - while neglecting the separate, real problem that your camera needs adjustment.

It would be wrong to interpret this as the object moving faster. The correct response to blur is to acknowledge that the image is blurred, describe the degree of blur as explicitly as the evidence allows, and govern accordingly: slow down, seek confirmation, and be explicit about the limits of current knowledge. Financial risk systems can make the wrong choice, especially in illiquid-asset contexts.

IV. Why the Confusion Matters

When financial models collapse epistemic uncertainty into a single volatility parameter alongside economic uncertainty, the resulting behaviour is subtle but pervasive.

Consider a mark-to-model asset revalued quarterly. Between valuation dates, the reported price does not move - what can be misinterpreted as economic stability is simply the absence of transactions forcing a repricing. When the revaluation occurs and the mark shifts materially, a naïve risk system reads the smoothed series as low-volatility until the revaluation cliff produces a spike, and then misreads the cliff as a sudden change in economic risk rather than the resolution of accumulated valuation uncertainty. The mark change may combine genuine economic movement, delayed recognition of prior shifts, and revised information quality. The volatility spike registered by the model conflates all three.

The practical consequences are visible in three places.

Hedge ratios become unreliable. The primary failure channel is correlation estimation from smoothed series. When a model uses stale prices to estimate the sensitivity of a private asset to a liquid benchmark, the resulting correlation is artificially low and the hedge ratio is liable to be understated. When revaluation arrives, the model observes a discontinuous move and may increase the hedge ratio aggressively, just as the primary market move may have already passed. The result may be a hedge ratio that is too low while valuation uncertainty is accumulating, followed by an aggressive recalibration once the delayed mark is recognised. The precise hedge error is case-specific, but the underlying problem is structural: stale or smoothed marks can make hedge calibration respond to measurement artefacts rather than economic exposure.

Volatility-sensitive frameworks can understate measured risk for appraisal-based or infrequently marked illiquid assets. Where stale or model-derived marks feed into volatility-sensitive internal models, portfolio optimisers, or VaR/Expected Shortfall-based allocation rules, the conflation can produce understatement of measured risk unless the framework adjusts for smoothing and valuation uncertainty. This issue is especially relevant for illiquid assets whose reported values depend on periodic valuation processes rather than continuous transaction-tested prices.

Hidden uncertainty accumulates. When epistemic uncertainty is invisible to the risk framework, it grows silently. Institutions do not routinely worry about what their models cannot see. Positions in opaque, illiquid instruments can expand without proportionate scrutiny. The gap between reported value and transaction-tested, observable exit or subsequently evidenced value widens until a liquidity event forces it closed. At that point, what had appeared to be a calm, low-volatility position reveals itself to have been, all along, a reservoir of unrecognised risk.

V. Where the Problem Becomes Most Visible: Illiquid Assets

Public equity markets are a special case. When a stock trades continuously on a liquid exchange, its price updates in near real-time against new information. The gap between reported price and

observable market value, while never zero, is generally tested more frequently by transaction and disagreement. Epistemic uncertainty has fewer opportunities to accumulate unnoticed; it does not persist as it does in illiquid markets.

Private assets are structurally different. A stake in an unlisted company, a senior secured loan to a mid-market borrower, a toll road concession, a portfolio of residential mortgages originated outside the agency market - all share a common feature: their value is often not continuously determined by trading. It is typically estimated periodically through valuation processes relying on models, assumptions, appraisals, comparables, broker indications or limited transaction evidence. All such processes embed judgement.

In these markets, valuation uncertainty is not continuously resolved. It accumulates between valuation dates, growing as the world moves while the model or valuation process stands still. A private equity portfolio may be marked using an EBITDA multiple calibrated to comparable public companies six months ago. A real estate asset may be appraised annually using comparables from a market that has since shifted materially. A structured credit instrument may be valued using a cash-flow model with default-correlation assumptions set at issuance and not refreshed with sufficient frequency. In each case the reported value carries epistemic uncertainty that is not reflected in any conventional volatility metric.

The consequence is the smoothing effect documented since Getmansky, Lo and Makarov (2004) and revisited in more recent work, including Coutts, Gonçalves and Rossi (2024): the apparent stability of illiquid-fund valuations through periods of public-market stress. Many appraisal-based or infrequently marked illiquid-fund series exhibit material positive serial autocorrelation - a statistical fingerprint of a smoothed series, in which each period's reported return is partly a delayed reflection of prior-period underlying economics.

The gap is most visible during market-stress episodes, where public markets and private markets do not reveal stress through the same measurement mechanism. Public-market indices reprice continuously through transaction-tested prices. Private-market indices often update through manager marks, appraisal cycles, fund reporting lags and valuation-process judgement. Table I below should therefore be read as a comparison of recognition mechanisms, not as a claim about relative economic performance.

This apparent stability is sometimes cited as evidence of lower volatility and superior risk-adjusted returns. That may be true for some assets and periods, but the table should not be read as proof of it. A safer inference: reported private-market series are often less frequently transaction-tested, more dependent on valuation process, and more exposed to reporting lag. Some of the observed stability may therefore reflect measurement architecture rather than economic calm.

Table I: How stress becomes visible under different valuation regimes

Episode	Public-market price discovery	Private-market valuation evidence	What the comparison shows
2008 crisis	S&P 500: approximately -37% on a 2008 calendar-year total-return basis; approximately -57% peak-to-trough across the 2007–09 cycle, depending on whether price or total-return series and daily or monthly observations are used.	Cambridge Associates US Buyout data, as analysed by Neuberger Berman, showed an approximately 28% peak-to-trough NAV decline during the 2007–09 global financial crisis.	Public markets showed continuous price stress; private-market stress was visible through reported NAV index drawdown and recognition cadence. The comparison is not a like-for-like economic loss measure. It illustrates that public and private markets reveal stress through different valuation and reporting mechanisms.
2022 repricing	S&P 500: approximately -18% in 2022; Nasdaq Composite: approximately -33% in 2022	Cambridge Associates US Private Equity Index: approximately -5.3% through the first half of 2022	Public markets reflected the rate shock with full-year index returns; private markets reflect first-half reported index performance. The key is the timing difference between continuous repricing and periodic valuation.

The discrepancy eventually resolves when liquidity events force private assets to market: a redemption or a fund’s end of life can each cause the accumulated valuation uncertainty to resolve in a single discontinuous event.

VI. Separating Economic Volatility from Valuation Uncertainty

Once the distinction is drawn, the appropriate response to each becomes correspondingly clearer.

Economic volatility is the territory of the existing toolkit: derivatives, dynamic hedging, diversification, and the probabilistic models of price evolution that underlie them. These work because economic uncertainty is fundamentally about the future trajectory of the world - a trajectory that, while unpredictable, follows statistical regularities that can be estimated, priced, and managed.

Epistemic uncertainty calls for something categorically different. The relevant tools measure and govern the quality of knowledge rather than the movement of prices. In practice that means:

1. explicit valuation uncertainty ranges around valuations rather than point-estimate NAVs;
2. explicit valuation uncertainty overlays or valuation reserves where appropriate, treated as governed judgement rather than modelled precision, and kept analytically separate from capital held against economic price movement;

3. valuation governance frameworks that distinguish between a valuation change driven by improved methodology and one driven by genuine economic movement; and
4. treating the staleness of data as a first-order risk factor rather than a footnote.

For clarity, this note uses internal economic capital to mean an institution's own capital allocation for risk management and decision-making purposes; valuation reserve to mean an adjustment or reserve against uncertainty in a valuation estimate; and regulatory capital only where a prudential regime specifically requires or recognises a capital deduction or requirement. Unless expressly stated, the argument here concerns internal economic capital and valuation reserves, not a claim about required prudential capital.

The practical point is to stop forcing both uncertainties into a single volatility parameter, and to maintain them as distinct quantities. Economic volatility governs price evolution, hedge ratios, and the dynamics of derivative pricing. Epistemic uncertainty governs the valuation uncertainty ranges around current estimates, any overlays or reserves applied to valuation uncertainty, and the governance requirements that apply to valuation processes.

The two interact, but they are not the same. An asset can have low economic volatility and high epistemic uncertainty: a long-dated infrastructure project whose revenues are contractually predictable but whose model relies on assumptions that have never been empirically validated. An asset can equally have high economic volatility and low epistemic uncertainty: a publicly traded commodity futures contract whose price moves dramatically but whose value is observed continuously and directly. Conflating them obscures both dimensions of risk.

Separating them does not eliminate valuation challenges in illiquid markets. It does allow those challenges to be clearly named and consistently measured rather than hidden inside a volatility statistic designed for an entirely different problem.

VII. Implications for Financial Markets

Two operational changes follow directly, with allocators as the primary addressee.

Reporting comes first. A disciplined framework would require illiquid books to carry explicit valuation uncertainty ranges alongside point-estimate NAVs, with valuation uncertainty overlays, reserves or internal economic capital treatment considered where appropriate. The analogy is to the logic of regimes such as EU prudent valuation, including the separate recognition of Market Price Uncertainty and Model Risk AVAs; it is not a claim that those prudential requirements apply directly to private holdings outside their legal perimeter. Existing supervisory, prudential and accounting regimes already recognise, in specific contexts, that observability, model dependence and valuation uncertainty require separate governance treatment. The missing step is less legal transplantation than conceptual discipline: applying the same distinction more consistently to private-market valuation infrastructure.

Hedge construction is the second. Hedge ratios derived from reported illiquid-asset returns should be stress-tested against desmoothed or proxy-based return estimates, rather than calibrated mechanically on stale reported marks. This approach draws on the desmoothing literature. Desmoothing is itself an estimator and not a perfect remedy: results are sensitive to specification, sample window, and the assumed lag structure, and desmoothed series should be treated as an input to risk-management stress tests rather than as a corrected ground truth. Using raw reported series for hedge calibration on illiquid books inherits the smoothing as a silent input, and the resulting ratios may oscillate against measurement artefacts rather than economic exposure.

Consider an illustrated case. A \$100m senior secured private credit position carries a reported quarterly volatility of 4% annualised, suggesting a 1-year 99% VaR of roughly \$9m. This is an illustrative normal VaR calculation using zero mean, annualised volatility and 99% confidence, and excluding default, loss-given-default, spread-migration and recovery-timing effects. Desmoothing the return series using a two-lag specification raises the implied volatility to approximately 7%¹, lifting the same VaR to roughly \$16m, and an explicit valuation uncertainty range of $\pm 5\%$ on the current mark identifies a further \$5m of valuation uncertainty that sits outside the volatility statistic entirely. These figures are illustrative only and are not a recommended calibration. Whether that uncertainty becomes an overlay, reserve, capital treatment or disclosure item depends on the institution's governance framework.

Allocators whose internal risk resources are mechanically linked to low reported volatility may be understating valuation uncertainty in private assets. Under the proposed framework, that uncertainty would be surfaced through ranges, overlays, reserves or internal capital treatment, depending on the institution's mandate and governance model. Where reported private-market returns are materially smoothed, methodologies that desmooth returns or separately recognise valuation uncertainty may reduce headline Sharpe ratios and change optimisation outputs. The scale and direction of the effect will depend on the asset class, dataset, smoothing specification, and allocation model used.

VIII. The Deeper Implication

Financial markets have historically pursued a particular ambition: eliminating uncertainty through models. If the world is uncertain, the response is a better model. If the model is uncertain, the response is a better calibration. The implicit goal is a system in which uncertainty is fully quantified, fully priced, and therefore fully tamed.

The ambition is not entirely misguided. Much uncertainty in financial markets can be productively modelled, priced, and distributed, and the achievements of quantitative finance in this domain are genuine and substantial. The ambition runs into a fundamental limit, however, when it encounters

¹ This example uses a simplified autoregressive desmoothing intuition. If reported returns are assumed to follow a smoothed process in which the observed return is a weighted average of current and lagged economic returns, desmoothing increases the estimated underlying volatility by reversing part of that smoothing. A two-lag specification is used illustratively here; the 4% to 7% uplift is not a universal conversion factor. It is intended to show the order of magnitude effect that can arise when stale or smoothed marks are used as inputs to volatility-sensitive risk measures. In practice, the specification, lag length and calibration should be independently tested.

epistemic uncertainty - the uncertainty that arises not from the unpredictability of the world but from the incompleteness of our knowledge of it.

Epistemic uncertainty is different, and resists each of the standard approaches. A better model does not eliminate it: the uncertainty is partly about the model itself. Reducing it to a single number discards information about its nature and source. It cannot be hedged cleanly in the same way as market beta. The available approach is explicit recognition: naming it, describing it as far as the evidence allows, measuring it where measurement is defensible, and treating it as a category of risk in its own right.

The deeper consequence is a different relationship with uncertainty than current practice has typically adopted. The defining task of risk governance becomes the separation of what can be priced from what can only be reserved against - and the honest acknowledgement that the second category does not vanish when ignored. Volatility, as conventionally measured, is a statistic about price evolution; it cannot also be expected to do the work of an estimate of valuation quality, and treating it as if it could is the source of the conflation described in this note.

Which uncertainty is being measured is the first question a serious risk framework has to answer.

Selected references and data sources

Board of Governors of the Federal Reserve System, Office of the Comptroller of the Currency, and Federal Deposit Insurance Corporation, SR 26-2: Revised Guidance on Model Risk Management, issued 17 April 2026; supersedes and replaces SR 11-7 and SR 21-8.

Board of Governors of the Federal Reserve System and Office of the Comptroller of the Currency, SR 11-7 / OCC 2011-12: Supervisory Guidance on Model Risk Management, issued 4 April 2011.

Regulation (EU) No 575/2013, Article 105; Commission Delegated Regulation (EU) 2016/101 of 26 October 2015 supplementing Regulation (EU) No 575/2013 with regard to regulatory technical standards for prudent valuation.

Basel Committee on Banking Supervision, Minimum capital requirements for market risk, including the FRTB Internal Models Approach and treatment of Non-Modellable Risk Factors.

International Accounting Standards Board, IFRS 13: Fair Value Measurement.

Financial Accounting Standards Board, ASC Topic 820: Fair Value Measurement.

Getmansky, Mila; Lo, Andrew W.; and Makarov, Igor, "An Econometric Model of Serial Correlation and Illiquidity in Hedge Fund Returns," *Journal of Financial Economics*, Vol. 74, Issue 3, 2004, pp. 529–609.

Couts, Spencer J.; Gonçalves, Andrei S.; and Rossi, Andrea, "Unsmoothing Returns of Illiquid Funds," *The Review of Financial Studies*, Vol. 37, Issue 7, 2024, pp. 2110–2155.

Cont, Rama, "Model Uncertainty and Its Impact on the Pricing of Derivative Instruments," *Mathematical Finance*, 2006.

Hansen, Lars Peter; and Sargent, Thomas J., "Robustness", Princeton University Press, 2008.

Glasserman, Paul; and Xu, Xingbo, "Robust Risk Measurement and Model Risk," *Quantitative Finance*, 2014.

Neuberger Berman, "The Historical Impact of Economic Downturns on Private Equity", December 2022, using Cambridge Associates LLC US Private Equity Buyout Index data. The paper reports an approximately 28% peak-to-trough NAV decline for US Buyout during the 2007–09 global financial crisis.

Cambridge Associates, *US PE/VC Benchmark Commentary: First Half 2022*, published 9 January 2023. Cambridge states that for the six months ended 30 June 2022, the Cambridge Associates LLC US Private Equity Index returned -5.3%, with quarterly returns of -0.4% in Q1 and -5.0% in Q2.

About this note

This is a research note published by Herfel og Rangen Ltd. Herfel og Rangen develops valuation infrastructure for private and illiquid instruments. Readers should treat this note as both research analysis and a statement of the problem space in which the firm operates. It reflects the author's analysis at the time of writing and is intended for professional investors, allocators, and practitioners working with illiquid assets.

References to regulatory regimes, supervisory guidance, and accounting standards are descriptive and conceptual. They are not legal or regulatory analysis and should not be read as a statement about how any specific institution is, or should be, supervised.

Nothing in this note constitutes investment advice or a recommendation to buy, sell or hold any security or to adopt any investment strategy. The note is for professional investors and is not intended for retail distribution in any jurisdiction.

No framework described here should be implemented without independent professional review appropriate to the institution's circumstances. The views expressed are the author's own and do not represent the position of any client, counterparty, or regulator.