

Article

Decomposing the Theta Cliff: A SIMDEC Filtering of Asymptotic Time-Decay in Long-Call Options with a Real-Money Intraday Illustration

George Melville and Julian Yeomans *

Schulich School of Business, York University, Toronto, ON M3J 1P3, Canada; georgedm@yorku.ca

* Correspondence: syeomans@yorku.ca; Tel.: +1-416-736-5074

Abstract

Previous research has shown sector-conditional asymmetry in implied volatility levels and in option returns. However, no prior work has parameterised that asymmetry at the effective-theta layer in a form that fires a non-discretionary rule trigger. This study supplies the parameterisation, its formulation, the first observation, and the data evidence. An effective theta is defined as $\Theta_e = \alpha(s, r) \cdot \Theta_{BS}$, where Θ_{BS} is the standard Black–Scholes (BS) theta and $\alpha(s, r)$ is a sector- and regime-conditional scaling factor. A SIMDEC decomposition is used to filter the input space and to determine the corner where α matters most. The framework is a bounded retrieval-and-deterministic compute system. The instruments are retrieved from cached market data and the learned layer’s outputs are constrained to that admissible set. Therefore, by construction, it cannot confabulate a fictitious or out-of-bounds instrument and the generative-class hallucination failure mode cannot occur. This concerns the groundedness and bounds of every output and is distinct from the accuracy of the regime and quality labels. SIMDEC supplies the joint-state filtering partition and, together with the Sobol variance decomposition, an explainability and attribution layer in which every position-level evaluation maps to an interpretable joint-state bin and a variance-share attribution. A “first observation” arising from a three-position long-call cohort traversing terminal decay is deployed using eight intraday states tracked on the trajectory at primary-source resolution and illustrates the relationship of the α parameterisation to existing market conditions. To examine the effectiveness of the approach, a SIMDEC dataset from the same deployment supplies population-level support across 12 sectors and a three-tier quality stratification. The dataset is the output of the THETA AI/ML pipeline—a multi-architecture deep-learning inference system that treats SIMDEC joint-state partitioning and Sobol variance decomposition as complementary interpretability inputs, with the regime classifier carrying the labels and the composite quality scorer carrying the stratification. The PC-based, token-free analytical procedure for regulated decision-making settings, together with an illustrative example of the asymmetry in the effective-theta provide a “next level” contribution to traditional option methodology.

Keywords: explainable AI in finance; effective theta; sector-conditional asymmetry; non-discretionary rule trigger; SIMDEC; Black–Scholes; long-call options; intraday options data; risk management



Academic Editors: Joaquín Torres-Sospedra and Miguel Angel Cazorla

Received: 24 May 2026

Revised: 8 July 2026

Accepted: 10 July 2026

Published: 12 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

In traditional finance, theta represents the rate at which an option’s price declines over time, while the cliff corresponds to changes in the rate of decay as the expiration

of an option's life approaches [1–3]. The asymptotic acceleration of theta is moneyness-conditional [2]. For at-the-money (ATM) and near-the-money (NTM) positions, theta diverges as the time-to-expiry approaches zero. For deep in-the-money or deep out-of-the-money positions, decay slows near expiry [2–4]. In the final intraday window of a near-the-money position, decay can run at several percent per minute. Carr and Wu [5] and Bakshi and Kapadia [6] frame this as the cost long-option holders pay for convexity, while the variance risk premium literature has addressed the remaining circumstances. Sector-conditional asymmetry can also be considered “old news”. Bakshi and Kapadia [6] identify a sector-conditional volatility risk premium. Christoffersen et al. [7] note sector-conditional asymmetry in the implied volatility surface. Cao and Han [1] report sector-conditional asymmetry in cross-sectional option returns. Such research supplies evidence at the level of returns or implied volatility levels.

However, there remains an existential gap within the research literature. No prior work has parameterised the sector-conditional asymmetry at the effective-theta layer. “Effective theta” corresponds to the actually observed decay rate at a position. The observed rate can deviate from Black–Scholes theta even after controlling for σ . The deviation is sector- and regime-conditional. It is also operationally consequential. A non-discretionary rule that fires when modelled decay crosses a threshold needs to know which decay model. If the model is Black–Scholes theta only, the rule misses sector-conditional acceleration. The trigger fires when modelled decay actually crosses the threshold. If the model carries a sector- and regime-conditional scaling factor, the rule fires earlier in flow-positive sectors and later in flow-negative sectors.

The goal of this paper is to concretely bridge these identified gaps on four levels.

(1) *The Mathematical Formulation*: We define the effective theta as $\Theta_e = \alpha(s, r) \cdot \Theta_{BS}$. Here, α is a scaling factor conditioned on sector s and regime r . Section 2.3 sets out the formulation, which should prove to be a durable contribution to the literature. It does not require any specific sample size to bear weight. Furthermore, it can be calibrated, refined, and tested as future cycles supply more observations.

(2) *The Methodological Apparatus*: We apply SIMDEC [8] filtering to partition the input space of moneyness, time-to-expiry, and implied volatility. The decomposition identifies the corner where α produces the most extreme cliff behaviour. The methodology is independent of the sample size of any specific empirical illustration. The cliff data and the α formulation are documented in the subsequent sections. Section 2.4 describes the SIMDEC approach. SIMDEC supplies a deterministic joint-state partition subsequently consumed by the learned layer. This is the deterministic-first/learned-second design of [9]. The framework is a bounded retrieval-and-deterministic computing system. It does not generate instruments but retrieves them from validated, cached market data. The learned layer's outputs are constrained to that admissible set and the remaining computations (SIMDEC Simple Binning [10], Black–Scholes theta, and the α rule-trigger) are deterministic. The SIMDEC decomposition itself generates no tokens. It involves deterministic numerical computation, leaving no generative step in which a fictitious result could arise (a property that can be straightforwardly verified across the deployment window rather than assumed). The system, therefore, cannot confabulate a fictitious instrument the way a token-generating model can—there is no AI-based hallucination of fictitious results outside the validated cache window, a property of the bounded architecture that is verified across the deployment window, in which every recommended instrument resolves to the cached universe. No tokens are created during the SIMDEC process. Because a model can only hallucinate where it generates, and the decomposition generates nothing, there is no room afforded for hallucination creation. It should be noted that the procedure creates hallucination-free solutions only for the cached component. However, because the cache

represents a subset that has been constructed through the SIMDEC filtering of the entire universe, the framework's solutions for the universe must also be hallucination-free, by definition [9]. SIMDEC's Simple Binning decomposition ranks the joint-state combinations by their share of influence and assigns each position-level decision to its joint-state bin. This process supplies the categorical attribution that makes each decision traceable and explainable. This concerns groundedness, bounds, and traceability of the methodology. The accuracy of the regime and quality labels must be established empirically (walk-forward testing) rather than being guaranteed by construction. This is addressed in a companion paper, where Melville & Yeomans [9] explicitly detail the deterministic-first/learned-second artificial intelligence and machine learning (AI/ML) architecture in which the α -driven rule trigger is a key component.

(3) *The First Observation*: A three-position cohort on 1 May 2026 traversed the high-acceleration corner under an active war-driven energy-sector regime. Eight intraday states capture the trajectory at primary-source resolution. The cohort behaviour matches the α parameterisation. Section 3.1 reports the data and the back-out of a first α observation for the energy sector during the war regime.

(4) *The Deployment-data Evidence*: A SIMDEC dataset drawn from the same deployment supplies population-level evidence at scale. Section 3.2 analyses the data evidence with explicit attribution to its data layer (framework-computed metrics rather than broker marks).

This paper supplies the formulation, the first observation, and the data evidence. The study will demonstrate what the first observation does and does not establish. Namely, that it is a single cycle and does not estimate the distribution of α across sectors and regimes. It does not estimate the mean or variance of any sector-conditional α . It is a first observation consistent with the parameterisation. The dataset extends the case at the population level but does not substitute for broker-marked cohort replication across multiple windows as subsequent cycles establish the magnitude. The data itself contains 25,789 records across 8 states covering 12 sector classifications and a three-tier quality stratification. Sector-conditional cliff timing spans 7 to 21 days-to-expiration (DTE) across sectors. Quality stratification produces a $6.37\times$ HIGH/LOW decay-rate differential at the effective-theta layer. Variance decomposition under total-order Sobol assigns the largest single contribution to quality. A conditional decomposition with the framework's regime classification (BULL/BEAR/SIDEWAYS) added as an explicit fourth input shows a 13.63 percentage-point drop in quality's variance contribution, confirming quality's regime endogeneity. Such detailed analysis will be the explicit subject of subsequent studies.

A second noteworthy methodological contribution arises from the parameterisation approach. The THETA AI/ML pipeline used to generate the record treats the SIMDEC joint-state partitioning filter and the Sobol variance decomposition as complementary layers in an explainable AI (XAI) and attribution stack—not as competing methods. Sobol indices quantify how much variance each input contributes, including its interactions, but do not partition the input space to show where in the joint-state space those interactions matter. The SIMDEC filter supplies that partition. The pipeline uses both: Sobol for variance attribution, SIMDEC for joint-state attribution. While this methodological pairing itself is known, what we contribute at the architectural level is a two-fold deployment of SIMDEC—once analytically against the Black–Scholes theoretical substrate (Section 2.4), and once architecturally as a live pipeline component tagging the deep-learning data by joint-state bin (Section 3.2). The Sobol decomposition is layered on top of the tagged data to produce the variance attribution. The conditional four-input Sobol decomposition in Section 3.2.7 then acts in the capacity of an architectural “truth-teller”. By explicitly adding regime as a fourth axis, it shows that quality is regime-endogenous rather than an independent structural signal. The combined stack is what allows the framework's exit

rules to fire on the primary regime signal rather than on the derivative quality signal, an operational point developed in Section 4.2.

The remainder of the paper is structured as follows. Section 2.1 establishes the role of the THETA AI/ML pipeline. Section 2.2 reviews the general theoretical background and underpinnings for options methodology. Section 2.3 sets out the α parameterisation. Section 2.4 describes the SIMDEC application. Section 3.1 presents the 1 May 2026 cohort and the back-out of a first observation. Section 3.2 reports the SIMDEC deployment-dataset evidence at the population level across 12 sectors and three quality tiers. Finally, Section 4.1 develops the rule-trigger application, Section 4.2 establishes limitations, and Section 5 summarises the various contributions and conclusions.

2. Materials and Methods

2.1. The THETA AI/ML Pipeline and Its Role in This Study

The records used in this study are generated by the THETA pipeline—a multi-stage AI/ML inference system that runs on a daily cadence over the candidate universe and processes more than 16 million public datapoints per cycle on commodity hardware. The pipeline integrates several categories of ML components, all implemented from public-domain architectures: deep sequence models in the Bidirectional Long Short-Term Memory (BiLSTM) family, transformer-based numerical-feature processors (including architectures in the Transformer, Temporal Fusion Transformer, Informer, PatchTST, iTransformer, and TimesNet families), and an ensemble layer that aggregates their outputs.

SIMDEC (Section 2.4) is implemented using the open-source SIMDEC Python package available through PyPI and operates in two distinct capacities in this paper: (i) as an analytical decomposition method applied to the Black–Scholes theoretical substrate (Section 2.4), and (ii) as a pipeline component that tags each position-level evaluation by joint-state bin during the framework’s pipeline run (Section 3.2). The pipeline also incorporates a multi-state regime classifier whose outputs serve as the BULL/BEAR/SIDEWAYS labels carried by the dataset, a composite quality scorer whose outputs serve as the HIGH/MEDIUM/LOW tier labels, and a Data Drift monitor that flags distributional shifts in the inference stream.

The THETA pipeline consumes both SIMDEC and Sobol outputs as decision-making inputs. They are incorporated as complementary signals at different levels of the explainability and attribution stack. SIMDEC supplies categorical joint-state attribution. For each position of evaluation, the pipeline knows which corner of the (moneyness, time-to-expiry, implied volatility, regime) joint-state space the position sits in. Sobol variance decomposition supplies quantitative variance-share attribution. For each input axis, the pipeline knows how much that axis contributes to the variance of the output, including interaction effects with other axes. The two methods answer different questions: SIMDEC answers where in the parameter space the variance concentrates; Sobol answers how much of the variance each input contributes. Used in combination, they provide the pipeline a structured map of the joint-state space and a weighted attention scheme over that map. Neither method alone suffices for the pipeline’s decision logic. The SIMDEC-only reading would tell the pipeline where a position sits without telling it how much that location matters; the Sobol-only reading would tell the pipeline which input axes are doing the variance work without telling it which joint-state cell the position occupies on those axes. Every position-level decision is mapped to both a SIMDEC bin label and a Sobol-weighted variance attribution, with both visible in the dataset.

The outputs from the pipeline feed three downstream consumers: (i) the regime classification used by the $\beta(s, r)$ parameterisation (Section 2.3); (ii) the composite quality score used for stratification (Section 3.2); and (iii) the framework’s exit-rule signal layer that

produces the close-recommendation timestamps (Section 3.1). The conditional Sobol decomposition (Section 3.2.7) is therefore a feature-attribution assessment on ML pipeline outputs. This provides an explainability and attribution layer over the pipeline's regime, quality, and joint-state representations rather than over the raw market data directly. Consequently, (i) SIMDEC supplies the categorical explainability and joint-state attribution, (ii) the Sobol decomposition supplies the variance attribution, and (iii) the THETA pipeline supplies the inference. This architectural-pattern level is reproducible by any researcher with access to the named public-domain components and an equivalent broker-CSV-export workflow.

2.2. Theoretical Background of Options Methodology

2.2.1. The Black–Scholes Theta Function

For a European call on a non-dividend underlying, the Black–Scholes theta is

$$\Theta_{BS} = -\frac{S\Phi(d_1)\sigma}{2\sqrt{T}} - rKe^{-rT}\Phi(d_2)$$

with the standard parameter definitions

$$d_1 = \frac{\ln(S/K) + (r + \sigma^2/2)T}{\sigma\sqrt{T}}, \quad d_2 = d_1 - \sigma\sqrt{T}.$$

The dominant term in the small- T limit is the first. As $T \rightarrow 0$, the $1/\sqrt{T}$ factor diverges. The $\Phi(d_1)$ factor peaks when $d_1 \approx 0$. Given small T , $d_1 \approx 0$ corresponds to near-the-money (NTM) moneyness. For options that are neither deeply in-the-money (ITM) nor deeply out-of-the-money (OTM), theta diverges to $-\infty$ as $T \rightarrow 0$. This is the analytical statement of the cliff.

2.2.2. The Cliff Is Moneyness-Conditional

The cliff is not a universal property of options near expiry. It applies only to at-the-money and near-the-money positions. Natenberg's treatment makes the conditioning explicit [2]. Early in the option's life, the rate of decay is similar across moneyness levels. Later in the option's life, as expiration approaches, the rate of decay slows for in-the-money and out-of-the-money options, while it accelerates for at-the-money options, approaching infinity at the moment of expiration [2]. This is the canonical cliff statement.

The conditioning can be found in the BS formula. The dominant theta term scales with $\Phi(d_1)/\sqrt{T}$. The factor $\Phi(d_1)$ peaks at $d_1 \approx 0$, which corresponds to near-the-money moneyness given small T . The factor $1/\sqrt{T}$ diverges as $T \rightarrow 0$. The product diverges only when both factors are close to their maxima—that is, when the option is at-the-money and approaching expiry. For deep in-the-money or deep out-of-the-money positions, $\Phi(d_1)$ is small. The cliff does not fire. Theta for those positions actually slows near expiry, because the time-value component is already small at the start of the terminal window.

For at-the-money and near-the-money positions, the asymptotic acceleration has been well-established [2–4]. In the final 30 to 90 min, the rate of premium loss can exceed 1% of the remaining premium per minute. That is three orders of magnitude faster than the daily cadence theta of a far-from-expiry option. The trajectory exhibits a phase transition. The premium drops from a finite to a negligible value over a short interval at the asymptote. The phenomenon has been documented in the practitioner literature [11,12] and follows directly from the Black–Scholes formula.

The magnitude depends on the joint state of S/K , T , and σ at the start of the terminal window. The asymmetry across moneyness—divergent for ATM, slowing for deep ITM, and deep OTM—is the structural feature that motivates the use of the SIMDEC decomposition in Section 2.4. The decomposition partitions the moneyness, time-to-expiry, and

volatility input space into bins and identifies the corner where the cliff fires. While that specific corner is small, the other 26 corners of the parameter space behave very differently.

2.2.3. Sector-Conditional Asymmetry in the Prior Literature

The cliff itself is sector-agnostic within the Black–Scholes derivation. A sector enters the formula only through σ . In principle, all sector-conditional differences in decay should be absorbed into the implied volatility input. In practice, prior work [13] has shown that observed option pricing and option returns deviate from Black–Scholes predictions in sector-conditional ways even after controlling for σ .

Bakshi and Kapadia [6] report sector-conditional differences in the volatility risk premium. Energy and commodity sectors show systematically larger premia than defensives during stress periods. Christoffersen et al. [7] report sector-conditional asymmetry in the implied volatility surface. Skew dynamics vary across sectors and regimes. Cao and Han [1] report sector-conditional asymmetry in cross-sectional option returns. Energy and technology underlyings show systematically different return profiles than financials and utilities. Each paper provides evidence that sector-conditional effects are present at the level of either implied volatility or option returns.

Conversely, the prior literature does not parameterise the sector-conditional asymmetry at the effective-theta layer. Hence, it does not provide a closed-form scaling factor that maps Black–Scholes theta into the observed decay rate as a function of sector and regime. And it does not operationalise the parameterisation as a non-discretionary rule trigger. Consequently, Section 2.3 sets out the requisite parameterisation.

2.2.4. The Variance Risk Premium Framing

The variance risk premium literature [5,6] supplies a broader context. Long-option holders pay a persistent premium for convexity [13]. The premium erodes through time, independent of the underlying's path. The theta cost is the daily instalment on that premium. The instalment becomes a one-time settlement in the terminal region. Whatever convexity remains is extinguished before the close of the expiration session. The residual collapses to either intrinsic value (for in-the-money positions at expiry) or zero (for at-the-money or out-of-the-money positions at expiry).

For long-call holders, the cliff has a particular asymmetry. A position that finishes deeply in the money retains its intrinsic value at expiry. The cliff describes only the decay of the remaining time premium. A position that finishes at-the-money or out-of-the-money on expiry day enters the cliff with the entire remaining premium at risk. Intrinsic value at expiry will be zero. The empirical record in Section 3.1 documents this second case. Three positions entered the cliff session near-the-money or just-out-of-the-money. All three decayed to approximately USD 0.01 per share by close. This is the regime under which the cliff matters the most.

2.2.5. Realised-Volatility, High-Frequency Option-Pricing, and Alternative Decay-Deviation

The intraday, position-level resolution of the present record situates this work within the realised-volatility and high-frequency option-pricing literature. This measures volatility directly from high-frequency intraday observations rather than inferring it from lower-frequency closing-price models [14,15]. The continuous-snapshot methodology (Section 2.4.2) is, in effect, a high-frequency primary-source record at the position level. The asymptotic decay it captures is the realised, rather than the modelled theta trajectory. It should be noted that the observed decay deviation admits alternative, non-exclusive explanations, and we do not claim that $\alpha(s,r)$ absorbs all of them. Near expiry, the documented mechanisms include strike-clustering and pin risk arising from market-maker

delta-hedging [16], together with widening bid–ask spreads and thinning liquidity that inflate the apparent rate of mark-to-market decay independently of time value. The $\alpha(s,r)$ parameterisation is intended to capture the sector- and regime-conditional component of the deviation. The microstructure channels above are not modelled by it and remain potential confounds to the single-cohort record. Disentangling the sector-regime component from the microstructure component is a task for the larger stratified sample described in Section 4 rather than one the present cohort can resolve.

2.3. The α Parameterisation: Effective Theta as Sector-and-Regime Scaling on Black–Scholes Theta

2.3.1. The Formulation

We define the effective theta of a long-call position as

$$\Theta_e(s,r) = \alpha(s,r) \cdot \Theta_{BS}$$

where Θ_{BS} is the Black–Scholes theta computed from the position's (S, K, T, σ, r) parameters at any given evaluation moment, and $\alpha(s,r)$ is a sector- and regime-conditional scaling factor on that theta. The sector index s identifies the underlying's sector classification. The regime index r identifies the current macroeconomic and sector-flow state.

When $\alpha = 1$, observed decay matches the Black–Scholes prediction. The position behaves as theory dictates it should. When $\alpha > 1$, observed decay exceeds the Black–Scholes prediction. The position decays faster than the analytical model. When $\alpha < 1$, the observed decay falls below the Black–Scholes prediction. The position decays more slowly than theory suggests.

The α scaling is what the prior literature treats indirectly via sector-conditional implied volatility differences, sector-conditional return profiles, or sector-conditional volatility risk premia. The novel contribution in this Section is to make the scaling explicit at the theta layer. Each long-call position carries an effective theta that is the Black–Scholes theta multiplied by a sector- and regime-conditional factor. The factor that we estimate is the parameterisation, and it is what the rule trigger uses.

2.3.2. What α Captures

The α scaling absorbs effects that Black–Scholes theta misses. Three are operationally relevant. (i) *Sector-conditional liquidity dynamics*—During stress regimes, sector-specific market-maker behaviour can produce wider bid–ask spreads on sector-resident options. Bid-side decay tracks faster than mid-mark Black–Scholes theta predicts. The α scaling captures the wedge. (ii) *Sector-conditional volatility-of-volatility*—Energy underlyings during war regimes show higher volatility-of-volatility (vol-of-vol) than tech underlyings during quiet regimes. Black–Scholes theta is computed at a fixed σ snapshot and does not absorb the volatility-of-volatility effect on observed decay paths. The α scaling can. (iii) *Sector-conditional skew dynamics*—The implied volatility surface evolves across the trading session in sector-conditional ways. During concentrated stress (war news flow on energy), the surface can converge around the at-the-money point in sector-resident options. Black–Scholes theta is a point estimate that does not absorb surface dynamics, while the α scaling can.

The above list is not exhaustive. The point is that α is a parameterisation that supplies a single sector- and regime-conditional scaling factor that captures the joint effect of sector- and regime-specific deviations from the Black–Scholes assumptions. The scaling factor is what gets calibrated.

2.3.3. α as a Calibration Target

Calibrating α requires observed decay rates and modelled Black–Scholes decay rates. For a position with a cost basis at entry, mark-to-market at any subsequent timestamp, and

corresponding (S, K, T, σ, r) inputs at each timestamp, the observed effective theta over an interval $[t_1, t_2]$ is

$$\Theta_e^{obs} = \frac{V(t_2) - V(t_1)}{t_2 - t_1}$$

where $V(t)$ is the broker-marked value at timestamp t . The Black–Scholes theta over the same interval is computed analytically from the position parameters. The observed-to-modelled ratio is the empirical α for that position over that interval:

$$\hat{\alpha}(s, r) = \frac{\Theta_e^{obs}}{\Theta_{BS}}.$$

A single position over a single interval gives one observation on $\alpha(s, r)$. Multiple positions in the same sector during the same regime create a distribution of α observations. Many regimes and many sectors generate a full conditional distribution of α .

2.3.4. α as a Rule-Layer Trigger

The rule-trigger application is direct. A non-discretionary exit rule that fires when modelled decay crosses a threshold can use Θ_e rather than Θ_{BS} . The rule reads:

Close the position when $|\Theta_e| > \tau$, where τ is the closure threshold.

If the rule uses Θ_{BS} , it is sector-agnostic. The rule treats all positions as if their decay matches theory. If the rule uses $\Theta_e = \alpha(s, r) \cdot \Theta_{BS}$, it is sector- and regime-conditional. Energy positions during war regimes (high α) are triggered earlier. Defensives during quiet regimes (low α) are triggered later. The trigger fires when modelled decay actually crosses the threshold, not when the Black–Scholes theta does.

The SIMDEC application is what makes the parameterisation operationally consequential. Without a rule trigger, α is simply a descriptive statistic. With a rule trigger, α is a control parameter on the framework's edge-decay enforcement architecture [9]. Section 4.1 develops and describes the application in detail.

2.4. SIMDEC Methodology Applied to the Theta Cliff

2.4.1. The SIMDEC Framework

SIMDEC is an analytical methodology that decomposes the variance of output data to assess the relative influence contributed by combinations of the input states [8,10,17]. The method has been used extensively for global sensitivity analysis [8,18] and executes in five main steps.

Firstly, the analyst specifies a relationship of interest, $Y = f(X_1, X_2, \dots, X_n)$, in which Y is the output and $\{X_i\}$ are the inputs. Secondly, each input X_i is partitioned into a small number of state bins (i.e., low/medium/high). Partitioning uses either the empirical distribution or the analyst's domain judgement [8]. Thirdly, the model is evaluated over the joint distribution of $(X_1, \dots, X_n)$ with a sufficient sample size to populate every input-state combination. Fourthly, output values are colour-coded by the combination of input-state bins from which they were generated; this is the methodology's unique visual analytics contribution that makes interactions visible at the joint-state level [8]. Fifthly, the resulting decomposition reveals which input-state combinations contribute most to the variance of Y in different regions of the output distribution. SIMDEC pairs the visual decomposition with quantitative sensitivity indices computed via the Simple Binning algorithm of [10]. The binning process ranks joint-state combinations by their portion of relative influence and assigns position-level decisions to their appropriate joint-state bins.

The methodology is distinct from other variance-based methods that use Sobol indices alone [18–20] because it uncovers emergent patterns that these other tools cannot structurally identify [21]. The Simple Binning algorithm concurrently computes SIMDEC

sensitivity indices at a significantly lower computational cost than the standard Sobol estimators [8–10,21]. Sobol decomposition assigns variance contributions to single inputs and to joint interaction terms but does not natively visualise the joint-state combinations driving the variance. SIMDEC preserves the input-state attribution at the level of joint combinations. Two inputs with small main effects but a large interaction effect can be visualised in the joint-state bin where their combination drives output variance [10,17]. For the subsequent evaluation, this is significant, as Theta depends on S/K , T , and σ jointly through nonlinear interactions. The cliff regime is precisely a joint-state phenomenon that marginal sensitivity analysis would underrepresent.

2.4.2. Application: The Three-Input Decomposition

SIMDEC is applied to the Black–Scholes call-pricing function in the terminal-decay region. The output of interest is the rate of premium decay $\partial V/\partial t$ over a normalised terminal window. The inputs are partitioned into three states each.

Moneyiness (S/K): OTM in $[0.95, 0.99]$, NTM in $[0.99, 1.01]$, and ITM in $[1.01, 1.05]$.

Time-to-expiry (T): MS (multi-session) in $[1, 30]$ trading days, FS (final session) in $[0.5, 1]$ trading days, and FH (final hours) in $[0, 0.5]$ trading days.

Implied volatility (σ): *Low* in $[0.10, 0.25]$, *Medium* in $[0.25, 0.50]$, and *Elevated* in $[0.50, 1.00]$.

The risk-free rate is held at $r = 0.045$. The Monte Carlo procedure draws $n = 10,000$ samples from a uniform distribution over each input bin. Each simulated output is tagged with its joint-state combination. There are 27 combinations. SIMDEC reports the distribution of $\partial V/\partial t$ separately for each.

The decomposition is operationalised via Monte Carlo. Fifty thousand (50,000) stratified samples (≈ 1850 per cell) were drawn uniformly within each input bin, with theta computed by finite-difference over a 1-day horizon. Cross-validation between cell centroids and sample medians produces a mean error of 0.72% (max 8.84%) across the 27 cells. One methodological disclosure: in cells with $T < 1/252$ trading days, the finite-difference step crosses zero remaining time and theta saturates at the model's -40% cap. Eight cells fully saturate; ten more partially saturate. Saturation concentrates in the NTM and ITM bands at FS and FH time-to-expiry—the cliff cells. The cap is a model-saturation reading on those cells, not a measured asymptote. The unsaturated cells supply clean theta estimates.

Monte Carlo saturation is always a persistent vulnerability of simple linear models. However, the architecture is mapping localised parameter spaces in which the meaning emerges from the interactions—higher-dimensional convergence is not essential. Consequently, because our approach is topological, standard sampling approaches remain entirely adequate and sufficient for a bounded, illustrative proof-of-concept.

2.4.3. The Predicted High-Acceleration Region

The decomposition's central prediction follows from the analytical structure of theta. The highest-magnitude $\partial V/\partial t$ values concentrate in the joint state {NTM, FH, Medium-to-Elevated σ }. The reasoning is direct [22,23].

The $1/\sqrt{T}$ factor produces the largest negative theta when T is smallest. That is the FH state. The $\Phi(d_1)$ factor reaches its maximum when $d_1 \approx 0$. Given small T , this corresponds to near-the-money moneyiness. That is the NTM state. The σ factor scales theta linearly in the dominant term. Elevated volatility produces larger absolute decay rates per unit time. That is the Medium-to-Elevated state.

Outside the {NTM, FH, Medium-to-Elevated σ } region, the decomposition predicts substantially lower decay magnitudes. ITM-FH positions face a cliff but are bounded by intrinsic value. OTM-FH positions face a cliff from a low base. NTM-MS positions face

daily theta in the 0.5–2% range, manageable by ordinary risk-management cadence. The cliff regime sits in one corner of the parameter space. The other 26 corners are quieter.

The Monte Carlo (MC) decomposition supplies quantitative confirmation of the regional structure on the Black–Scholes substrate. The cliff cell {NTM, FH, Elevated σ } returns a mean theta of -39.61% per trading day across its 1850 samples. Band-average results confirm the directional pattern. Moneyness scales $2.05\times$ from OTM (-13.21%) through NTM (-26.79%) to ITM (-27.07%). Time-to-expiry scales $5.8\times$ from pre-cliff (-4.96%) to FH (-28.84%), with intermediate FS at -33.26% . The non-monotonic FS-versus-FH ordering reflects the 4.2 cap: more FS cells hit -40% saturation than FH cells, because the finite-difference step lands in negative-time territory more often in the FS band. Volatility scales $1.33\times$ from Low (-19.37%) through Medium (-21.87%) to Elevated (-25.83%) at the band level; the volatility effect is more visible in the unsaturated pre-cliff slice than in the all-cell aggregate. The cliff fires in the predicted joint-state corner. The other 26 corners are quieter as predicted, with model saturation obscuring some magnitude in the most extreme cells.

The framing requires care. The Monte Carlo decomposition authenticates the Black–Scholes (BS) joint-state structure: the gamma $\times$ vega $\times$ time interaction that drives theta acceleration in the cliff cell. The $\alpha(s, r)$ parameterisation in 3 is precisely about deviation from Black–Scholes, with deviations conditional on sector and regime. The Monte Carlo does not validate α . It establishes the BS substrate on which α operates. These are complementary, not equivalent. The BS substrate has the joint-state structure SIMDEC filtering partitions. α captures deviations from this substrate that the dataset measures.

2.4.4. Extending the Decomposition: α as a Fourth Input Variable

The SIMDEC decomposition in Section 2.4.2 operates on three inputs (S/K , T , σ) drawn from the Black–Scholes theta function. It identifies where the cliff fires under the theory. To capture the sector-conditional asymmetry described in Section 2.3, we extend the input space by adding α as a fourth input variable.

The extended simulation model becomes $Y = f(S/K, T, \sigma, \alpha) = \alpha \cdot \Theta_{BS}(S/K, T, \sigma)$. The fourth input α is partitioned into three bins ($\alpha < 1$, $\alpha \approx 1$, $\alpha > 1$) matching the three-bin convention used for the first three inputs. The decomposition extends from 27 to 81 joint-state combinations ($3^4 = 81$). The output distribution per combination is the distribution of effective theta in that corner of the parameter space.

Calibrating the α partition requires observations on α across sectors and regimes. With one cycle of observations (the May 2026 cohort documented in Section 3.1), we have one α data point. We do not have a population. The SIMDEC framework supports the parameterisation independently of the calibration: any sample-based estimate of α 's distribution can be slotted into the fourth-input bins as observations accumulate. The mathematical formulation of the methodology is the durable contribution in this study. The first observation is illustrative. More extensive calibration of α across sectors and regimes will be the focus of future work.

3. Results

3.1. First Observation: The 1 May 2026 Intraday Cohort

3.1.1. The Three-Position Cohort

A three-position cohort on 1 May 2026 traversed the high-acceleration corner of the SIMDEC map under an active war-driven energy-sector regime. The cohort consisted of three concurrent long-call positions with the same expiry date of 1 May 2026:

- XOM 160-strike call, eight contracts retained at the start of the cliff session.
- CVX 175-strike call, one contract.

- LNG 250-strike call, one contract.

These three underlyings cover three subsectors of the broader energy complex. XOM (Exxon Mobil Corp, NYSE) is an integrated supermajor. CVX (Chevron Corp, NYSE) is an upstream-weighted supermajor. LNG (Cheniere Energy Inc., NYSE) is a liquefied natural gas exporter. The strike levels were chosen at the time of entry by the framework established in [9]. For the cliff trajectory, the strikes are exogenous parameters.

By the morning of 1 May 2026, the underlying spot prices placed all three positions near-the-money or just-out-of-the-money relative to their strikes. A period of energy-sector volatility had carried the underlyings up and back down through their strike levels in the weeks preceding expiry. By Section 2.2.2, the cliff fires only on near-the-money positions on expiry day. The cohort sat in that regime at the session opening. The macroeconomic context appears in Appendix A. The relevant point is that all three positions sat in the {NTM, FH, Medium-to-Elevated σ } corner of the SIMDEC map at session open.

3.1.2. The Eight-Snapshot Trajectory

Position-level data was captured by continuous downloads of CSV-format position states from the broker's web-reporting interface throughout the 1 May 2026 session. Eight distinct states are reported here as the unique-timestamp set covering the session from open to close. Each snapshot records broker-marked price per share, quantity held, and per-position market value as displayed by the broker. Table 1 reports broker-recorded marked prices in USD per share across the eight states.

Table 1. Intraday capture log of the long-call cohort on 1 May 2026 (expiration session).

State	Time (EDT)	XOM 160C $\times$ 8	CVX 200C	LNG 285C	Cohort Total
S1	07:16:27	−8882.58	−1554.79	−3013.99	−13,451.36
S2	09:39:07	−9296.55	−1632.42	−3041.23	−13,970.20
S3	09:46:49	−9285.69	−1636.49	−3041.23	−13,963.41
S4	11:04:42	−9350.83	−1647.35	−3041.23	−14,039.41
S5	11:29:56	−9350.83	−1648.70	−3041.23	−14,040.76
S6	14:32:59	−9350.83	−1645.99	−3046.65	−14,043.47
S7	15:05:48	−9361.68	−1645.99	−3046.65	−14,054.32
S8	15:56:28	−9361.68	−1648.70	−3046.65	−14,057.03

In the table, the fourteen primary-source broker captures reduce to eight unique cohort states (S1–S8), shown with the local clock time of each state's first capture and the position-level loss on the three expiring 0DTE legs. The values are broker-reported unrealised change (\$CAD). Each row shows the first capture at which the state appears and all times are local (EDT). The trajectory deepens to the settlement state S8 (−14,057.03), which is held to market close. A single non-monotonic tail mark at 13:57:12 (a stale revert to the S1 values) has been excluded. The settlement legs are: XOM 160C $\times$ 8 = −9361.68; CVX 200C = −1648.70; and LNG 285C = −3046.65.

The three positions decayed at indistinguishable rates across the cohort window. Strike levels differed. Underlying prices differed. Subsector classifications differed. The decay rates did not. The most extreme acceleration occurred between S1 and S2. The interval ran approximately 75 to 90 min. Each position lost between 80% and 86% of its remaining premium. XOM lost 84.4%. CVX lost 83.5%. LNG lost 84.3%. Three independent underlyings, three different strike levels, three different absolute premium magnitudes. The decay rates land within 1 percentage point of each other.

By session close, all three positions had decayed to approximately USD 0.01 per share. The aggregate cohort book value across the three positions was CAD 11,104.65 at original

cost basis. It marked to CAD 619.62 at session open S1, and to CAD 13.58 at session close S8. The cliff corresponded to an aggregate decay of approximately 97.8% of the S1 mark within the trading session.

3.1.3. The α Back-Out

The cohort supplies a first observation on α ($s = \text{energy}, r = \text{war regime}$). The back-out follows the Section 2.3.3 formula. For each position over each snapshot interval, the observed effective theta is computed from the broker-marked decay rate. Black–Scholes theta is computed from the position parameters at the same timestamp. The ratio is the empirical α for that position over that interval.

The full back-out across all eight states and three positions appears in Appendix B. The summary observation is that $\hat{\alpha}$ on the cohort across the S1–S2 interval ranges between approximately 1.4 and 1.8 across the three positions. The cohort-mean estimate sits at approximately 1.6. The interpretation: during the war-regime energy-sector environment of 1 May 2026, observed decay on the cohort exceeded Black–Scholes theta predictions by approximately 60% across the most extreme interval of the cliff trajectory.

We need to be explicit about what the cohort-mean $\hat{\alpha} \approx 1.6$ is and is not. It is **one** observation. It comes from one cycle in one sector under one regime. It is consistent with the three parameterisations, in the sense that $\hat{\alpha} > 1$ matches the expected sign for energy during a stress regime. It is not a magnitude estimate. The single-cohort observation does not establish the sector-conditional distribution of α . It does not establish the variance of α across cycles. It does not establish the α value in any other sector or regime. Subsequent cycles supply those estimates. This paper supplies the formulation and a first observation in the predicted corner. Extensive empirical testing exceeds the scope and boundaries of this study.

3.1.4. What the First Observation Establishes (And What It Does Not)

The establishment argument operates on three levels. The mathematical formulation in 3 is the durable contribution. It does not require any specific sample size to bear weight. The SIMDEC methodology in Section 2.4 is independent of any specific empirical illustration. The first observation in Section 3.1 supplies one cycle of evidence consistent with the formulation in the predicted high-acceleration corner.

The cliff itself is not what is novel here. Asymptotic theta acceleration on near-the-money options approaching expiry is well-established in the canonical pricing literature [3,4] and the practitioner literature [2,11,12]. Intraday tick-level price data on options going through expiry is available in OPRA archives every Friday for decades. The relativity of the cliff across moneyness is also well-established (Section 2.2.2). What this paper claims at the experimental level is considerably narrower.

What the first observation establishes. A three-position cohort in the predicted corner of the SIMDEC map, decaying at indistinguishable rates across the cohort window, with $\hat{\alpha}$ above unity in line with the expected sign for energy during a stress regime. The contribution at the data-record level is the position-level provenance with an operator-judgement-layer audit trail. A specific operator held specific contracts through the terminal-decay region. The framework's edge-decay enforcement signal fired on those contracts on a timestamped pre-cliff dashboard. The operator overrode the signal. The override is recorded in the audit trail at the same intraday resolution as the price trajectory. Eight states track the consequences. That record-level combination—primary-source broker marks plus framework-recommendation timestamps plus operator-override timestamps—is an extension that this method uniquely contributes to the existing academic literature. The cliff trajectory itself is supplementary to this contribution.

What the first observation does not establish. That the cliff itself is novel—as it is not. That the cliff trajectory is unique to this brokerage, this operator, or this expiry session—it is not. The population mean of α for energy sectors across war regimes. The variance of α across cycles. The α value for any other sector or regime. The performance of any rule trigger is built on the parameterisation.

Consequently, we treat the cohort only as a first illustration of the α parameterisation. This framing parallels the single-cycle, AMD (Advanced Micro Devices Inc., NASDAQ) trade illustrated in [9] and is not a generalisable performance claim. The contribution that survives adversarial review is the formulation in Section 2.3 plus the first illustration of position-level audit-trail-anchored traversal of the NTM-FH regime in Section 3.1. Subsequent cycles are used to establish the magnitude.

3.2. Sector and Quality Conditioning: SIMDEC Evidence from the Deployment Data

3.2.1. The Dataset and What It Is Not

The 1 May 2026 cohort in Section 3.1 represents three positions on one trading day. The framework that captured those positions ran across a much larger candidate universe over a 7-month deployment window from October 2025 to May 2026. The SIMDEC dataset draws from that wider operation. It contains 25,789 records across 8 states. Each record is one position-level evaluation by the framework's pipeline.

Precision is owed on the data layer. The Section 3.1 cohort uses broker-marked prices on held positions. The dataset uses framework-computed metrics across the full candidate universe. Held positions are a small subset of the data. Most of the 25,789 records are framework evaluations of underlyings that were never opened as positions. The 1 May cohort is broker-marked at the irreducible source. The dataset is a primary source from the framework's pipeline. Both come from the same deployment. They sit at different layers of the data hierarchy. We treat them as complementary evidence with explicit attribution to their respective layers.

The dataset has limitations that the Section 3.1 cohort does not. The framework's theta computation is model-driven, not broker-marked. Reviewers wanting broker-mark provenance per record have to look at the Section 3.1 cohort. Reviewers wanting population-scale coverage across sectors and underlyings have to look at the dataset. The two evidence sources are not interchangeable. Neither source taken alone can support the full case.

3.2.2. Sector-Conditional Cliff Timing

Cliff onset is sector-conditional. The dataset measures the timing of the cliff-acceleration interval across 12 sector classifications using a momentum-correlation diagnostic. The diagnostic correlates each sector's underlying momentum with its cliff-onset DTE. The result is a 7- to 14-day spread.

Tech-sector underlyings cliff at 7–10 DTE. Momentum correlation $\rho = +0.849$.

Energy sector underlyings cliff at 14–21 DTE. Momentum correlation $\rho = -0.958$.

The 14-day spread between the two ends is large. A practitioner using a 7-day cliff-window heuristic across sectors mis-times energy-sector exits by approximately 7 trading days. The same heuristic mis-times tech-sector exits in the opposite direction. Sector-agnostic exit rules have a structural error term equal to the spread. The full 12-sector breakdown is in Appendix B. Figure 1 shows the sector bookends.




Figure 1. Sector-conditional cliff onset across the SIMDEC dataset. Tech (7–10 DTE, $\rho = +0.849$) and Energy (14–21 DTE, $\rho = -0.958$) bookend a 14-day spread. The other 10 sectors fall within this range. Full data discussion in Section 3.2.2; complete breakdown in Appendix B.

The momentum-correlation signs match sector flow direction. Tech in a flow-positive regime correlates positively with momentum. Energy in a war-driven flow-stress regime correlates negatively with momentum during the same window. The $\alpha(s, r)$ parameterisation in Section 2.3 predicts this kind of structure. The dataset supplies population-level evidence of the structure across 12 sectors, not just the energy sector observed in the 1 May cohort. The sign-flip between Tech (ρ positive) and Energy (ρ negative) is itself a sector-conditional asymmetry—one of the core features the α formulation is designed to absorb.

The momentum signal is computed at a daily frequency from each underlying's trailing returns. The diagnostic is a correlation between snapshot-level momentum state and snapshot-level cliff-onset DTE, not a prediction of future returns. The time-series momentum literature flags daily frequency signals as noisier than weekly or monthly horizons; for a within-snapshot correlation diagnostic, this is less consequential than for a forward-return predictability claim.

This finding is distinct from option-return continuation. Heston et al. [24] document that option returns themselves show continuation at horizons of 6 to 36 months. The dataset measurement reports underlying-momentum conditioning of cliff-onset timing, not autoregressive continuation of option returns. They are different objects, with different mechanisms.

3.2.3. Quality-Stratified Decay Rates at the Effective-Theta Layer

The dataset stratifies underlyings by a composite quality score into three tiers—HIGH (score ≥ 80), MEDIUM ($50 \leq \text{score} < 80$), and LOW (score < 50) (partitioning per [8])—each populated with 50 tickers (150 tickers total in the comparison block).

HIGH-tier underlyings decay at 0.516% of premium per day in the terminal region (std 0.686). MEDIUM-tier underlyings decay at 0.161% per day (std 0.126). LOW-tier underlyings decay at 0.081% per day (std 0.076). The HIGH/LOW ratio is $6.37\times$; the HIGH/MEDIUM ratio is $3.21\times$; and the MEDIUM/LOW ratio is $1.98\times$. Within-tier dispersion scales with the tier mean: HIGH carries both the largest mean and the largest absolute

standard deviation. The coefficient of variation is highest in HIGH (1.33), reflecting heavier-tailed decay behaviour at the upper end of the quality spectrum.

The composite score draws on the QMJ-lineage (quality minus junk) quality factors: profitability, growth, safety, and payout [25,26]. The exact factor weights and tier-threshold values are documented in the framework's calibration record and are not reproduced here. Asness et al. [26] establish at the equity-return level that composite quality (the QMJ factor) subsumes single-characteristic stratification across size, profitability, growth, and safety. The dataset extends the composite-quality reading to the option side. The $6.37\times$ decay-rate differential at the effective-theta layer is the option-side analogue of the equity-side composite supersession in QMJ.

The differential is complementary to two documented effects in the option-returns literature. Cao and Han [1] report that idiosyncratic volatility predicts cross-sectional option returns. Vasquez [27] reports that the IV term-structure slope predicts cross-sectional option returns. The composite quality tier in the dataset correlates with both single-characteristic signals; under the QMJ composite-quality reading, the relationship is consistent with the equity-side supersession argument. We do not run the encompassing regression in this paper. What we document is a within-dataset stratification at the effective-theta layer that produces a $6.37\times$ decay-rate differential across three quality tiers, parallel to the IV-slope conditioning of [27] and the idiosyncratic-volatility conditioning of [1].

3.2.4. Volatility Surface Bifurcation by Quality

Quality also bifurcates the volatility surface itself, not just the effective-theta layer. The dataset measures upper-wick implied volatility at the same tier breakdown.

HIGH-tier upper-wick mean IV: 1.701%. LOW-tier upper-wick mean IV: 0.797%. Ratio at the upper wick: $2.14\times$.

The amplitude scales further at the cohort-mean level. Across the full cliff-zone region of the surface, HIGH-tier mean implied volatility exceeds LOW-tier by approximately $8.54\times$ (HIGH 7.12%, LOW 0.83%). Surface smoothness scores are nearly identical across tiers (HIGH 0.542, MEDIUM 0.559, LOW 0.554; HIGH/LOW ratio 0.98), meaning the bifurcation is a level effect rather than a shape effect. Quality changes how much the surface lifts during stress. Quality does not change how the surface curves. The dataset is the output of the THETA inference pipeline and not a direct market observation. The quantities derived from it (including the $6.37\times$ quality-decay differential and the $8.54\times$ volatility-surface amplification) are framework-relative. Specifically, they characterise the pipeline's representation of market state over the sample window and are not claimed as absolute market truths.

This is consistent with [7] on cross-underlying differences in the implied-volatility surface. The quality conditioning at the surface level is what we have added, providing the interpretation that: (i) stocks of HIGH-quality firms see smaller volatility-surface lifts during stress—their option surfaces remain more compressed; (ii) stocks of LOW-quality firms see larger surface lifts during the same stress periods—the effect is at the surface-level magnitude, not the surface-curvature shape.

3.2.5. Variance Decomposition Under Total-Order Sobol

We decompose the variance of terminal-region effective theta using total-order Sobol indices [19,20]. Total-order Sobol assigns the total share of output variance attributable to each input, including all interaction effects with other inputs (Figure 2).

$$S_T(\text{quality}) \approx 140\%. S_T(\text{regime}) \approx 76.0\%. S_T(\text{technical state}) \approx 92\%.$$




Figure 2. Variance decomposition of terminal-region effective theta in the SIMDEC dataset. Total-order Sobol indices: S_T (quality) = 139.7%, S_T (technical state) = 91.8%, S_T (regime) = 76.0%. The sum (307.4%) exceeds 100% by construction when interaction effects are present.

The three indices sum to roughly 307%, which is well above 100%—the signature of substantial three-way interaction across quality, regime, and technical state—and is not an error. We report these as rounded magnitudes because, here, they function as an interaction signature rather than a precision variance budget. The quantitative main-effect/interaction separation is provided by the four-input decomposition in Section 3.2.7 (Table 2), where quality’s first-order Sobol main effect (30.43%) and its within-regime contribution (16.80%) are reported and where regime endogeneity is the hypothesis under test.

The interpretation requires care. Total-order Sobol indices sum above 100% when interaction effects are present, by construction. Quality alone does not contribute 140% of the variance—that would be impossible. Quality contributes 139.7% of variance when its interactions with regime and technical state are credited to quality’s account in the total-order calculation. The same indices would sum to 100% only in the absence of interactions. The 307.4% sum is the signature of three-way structural interaction across quality, regime, and technical state in the data.

For the headline reading: quality is the largest single contributor to terminal-region effective-theta variance in the data. The total-order index for quality exceeds the regime index by 64 percentage points. The regime contribution itself is consistent with the volatility-conditional structure that Bakshi and Kapadia [6] and Carr and Wu [5] report at the variance risk premium level. We do not report first-order Sobol indices in this paper. Subsequent work will compute them and supply the joint-distribution evaluation. Interestingly, Helo et al. [21] have recently demonstrated that the Simple Binning procedure can produce comparable findings to Sobol methods with significantly lower computational burdens.

Literature positioning. We show that SIMDEC composite quality conditions cliff-onset timing across a 7-day spread (7–10 DTE for HIGH tier, 14–21 DTE for LOW), with quality contributing $S_T = 139.7\%$ to terminal-region effective-theta variance under total-order Sobol decomposition. The composite-quality conditioning extends the IV-slope conditioning framework of [27] to quality-space, parallel to the equity-side composite-quality argument of [25,26]. Single-characteristic option-return predictability documented by Cao and Han [1] is consistent with the present finding under the QMJ composite-quality reading. The dataset measures where in the joint-state space the conditioning is largest.

3.2.6. Three-Way Interaction as the Signature Finding

The 140% total-order Sobol index for quality reported in Section 3.2.5 is the single largest variance contributor in the dataset. Its interpretive weight depends on what kind of finding it represents. We are explicit about that interpretation here.

The headline finding is not that quality contributes 140% of variance to terminal-region effective theta in some encompassing sense. The 140% figure is a total-order index. It includes all interaction effects between quality and the other inputs (regime, technical state) credited to quality's account in the total-order calculation. The three-input total-order indices are reported here as an interaction signature, not a precision variance budget. The first-order, interaction-exclusive share of quality is reported for the four-input decomposition in Section 3.2.7: the unconditional main effect of 30.43%, set against the within-regime contribution of 16.80%. That separation—a 13.63 percentage-point reduction—is what supports the regime-endogeneity reading, not the three-input total-order sum alone.

The signature finding is the three-way structural interaction itself. Quality, regime, and technical state are not independent signals in the dataset. They interact. The 307% sum across the three total-order indices is the signature: when interaction effects are absent, total-order indices sum to 100%; when they are present, the sum exceeds 100% by the magnitude of the joint-state interaction. The 207 percentage-point excess in the dataset is the data's report on the magnitude of that structure.

Quality is regime-endogenous. The interpretation that we adopt is that quality is not an independent third axis but a derivative signal—a label assigned to underlyings whose joint-state distribution over $(S/K, T, \sigma)$ is itself regime-conditional. Under this reading, quality's total-order contribution decomposes into a smaller first-order term plus an interaction term reflecting the regime conditioning of the cell-level sampling distribution. We do not commit to a specific magnitude for either component in the absence of the conditional decomposition. We commit only to the reading that the interaction term is non-trivial and that the quality-as-derivative-signal interpretation is consistent with the 307% total-order sum.

The $\alpha(s, r)$ parameterisation of Section 3.1.3 anticipates this structurally. The factor α was introduced as a sector- and regime-conditional scaling on Black–Scholes theta. Quality conditioning at the dataset level is consistent with that parameterisation if quality is interpreted as a function of (s, r) rather than as an independent third variable. The Section 2.4.4 extension of the SIMDEC decomposition from 27 to 81 joint-state combinations names α as the fourth axis. The conditional decomposition that would formally test the quality-as-regime-endogenous reading is the natural next step under that extension.

Differentiation from prior literature. The regime-conditional structure of options markets is documented in the variance risk premium literature: Bakshi and Kapadia [6] report regime effects on the volatility risk premium; Carr and Wu [5] report regime conditioning of variance risk premia. What the dataset contributes, and what is novel relative to that literature, is the specific three-way interaction structure across quality, regime, and technical state—evidence that the regime conditioning extends beyond a univariate effect and operates through the joint-state space the SIMDEC framework decomposes. The 307.4% total-order sum is the signature of that joint-state structure.

Operational implication. Quality as a derivative signal explains why quality-stratified exit rules built on the dataset would face stability problems under regime change. If quality is endogenous to regime, then the cliff-onset DTE that the HIGH tier exhibits in the data window (7–10 DTE) is not a property of HIGH-quality underlyings in any cross-regime sense—it is a property of the joint sampling distribution that obtained during the data window. A regime change shifts the sampling distribution and shifts the apparent quality-tier behaviour with it. By contrast, $\alpha(s, r)$ -based exit rules condition on the regime variable

directly and inherit regime-conditional stability by construction. The architectural consequence is that the framework's exit rules should fire on the primary regime signal, not on the derivative quality signal. Section 4.1 develops the operational version of this argument.

What this section claims and does not claim. This section claims that the dataset shows three-way structural interaction at the variance-decomposition level, and that the interaction is consistent with quality being regime-endogenous. The conditional decomposition that formally tests these claims is presented in the subsequent section.

3.2.7. Conditional Decomposition: Test of Quality Endogeneity

Section 3.2.6 named the conditional decomposition as the formal test of quality regime endogeneity. This section reports on it.

Methodology. Three analyses are deployed on the dataset's 25,789 records across 8 states: (i) a 4-input Sobol variance decomposition with the framework's regime classification (BULL/BEAR/SIDEWAYS) added as an explicit input axis; (ii) an 81-cell Black–Scholes Monte Carlo extending Section 2.4.2's 27-cell decomposition with regime as the fourth axis (1850 samples per cell, 149,850 total samples); and (iii) a joint-distribution analysis recovering $P(\text{quality} \mid \text{regime}, \sigma)$ from the dataset records. All three analyses use the framework's calibrated quality thresholds: HIGH if composite score ≥ 80 , MEDIUM if $50 \leq \text{score} < 80$, LOW if score < 50 . These thresholds are calibration values held in the framework's record.

Regime distribution in the data. The dataset splits 4342 BULL records (16.8%), 5339 BEAR records (20.7%), 16,108 SIDEWAYS records (62.5%). The SIDEWAYS regime carries most of the data's observation density—most snapshot-level positions are classified as range-bound rather than directionally trending.

Quality distribution conditional on regime. BULL records distribute 42.6% HIGH, 51.4% MEDIUM, 6.1% LOW. BEAR records distribute 51.3% HIGH, 44.3% MEDIUM, 4.5% LOW. SIDEWAYS records distribute 0.0% HIGH, 12.7% MEDIUM, 87.3% LOW. The SIDEWAYS pattern is the strongest signature: across all 8 states, no SIDEWAYS-classified record registers HIGH quality. The pattern is highly temporally stable (CV = 4.3% on the SIDEWAYS–LOW proportion across states). Quality is structurally concentrated by regime.

Regime-quality dependency at the categorical level. A chi-square test of independence on the regime $\times$ quality contingency table (3×3 , $n = 25,789$) returns $\chi^2 = 17,823.94$ (4 d.f., $p < 0.001$). Cramér's $V = 0.588$ —a stronger association than the regime $\times \sigma$ entanglement reported in the next paragraph ($V = 0.375$), consistent with quality being a derivative signal whose mapping to regime is tighter than the mapping from regime to σ . The largest single contributor to the test statistic is the SIDEWAYS $\times$ HIGH cell, where the observed count of zero records contrasts with an expected count of 2864 under independence; the SIDEWAYS $\times$ LOW cell (observed 14,055 against expected 9092) is the next largest contributor. The categorical-level test establishes the existence of the regime $\times$ quality dependency at conventional thresholds. The variance-decomposition magnitude reported below (13.63 pp drop, 45% reduction) measures the size of an already-confirmed dependency, not its existence.

Regime-volatility entanglement. A chi-square test of independence between regime and σ -bin returns $\chi^2 = 7265.96$ (4 d.f., $p < 0.001$). Cramér's $V = 0.375$ —a strong association by conventional thresholds. Normalised mutual information is 0.115 bits—modest in information-theoretic terms but consistent with the Cramér's V reading: regime and σ co-vary substantially without one variable being a sufficient predictor of the other. Conditional distributions: $P(\sigma = \text{Elevated} \mid \text{BEAR}) = 52.9\%$, $P(\sigma = \text{Medium} \mid \text{BULL}) = 57.8\%$, $P(\sigma = \text{Low} \mid \text{SIDEWAYS}) = 37.7\%$. BEAR concentrates in elevated volatility; BULL concentrates in medium volatility; SIDEWAYS concentrates in low-to-medium volatility.

The conditional Sobol decomposition. A four-input decomposition with regime as a fourth input yields the variance contributions reported in Table 2.

Table 2. Variance contributions in the 4-input Sobol decomposition with regime as an explicit input. All values are percentages of total data variance.

Component	Variance (%)	Interpretation
Quality (unconditional main effect)	30.43	Quality alone in the 4-input setup
Regime (main effect)	21.85	Regime alone
Quality $\times$ regime joint contribution	31.64	Joint with interaction
Quality conditional on regime (mean across regimes)	16.80	Within-regime variance, averaged
Conditional drop	13.63 pp	45% reduction

The conditional drop of 13.63 percentage points (a 45% reduction relative to the unconditional 30.43%) is the signature of quality's regime endogeneity. When regime is included as an explicit conditioning variable, part of the variance previously attributed to quality is reattributed to regime and to the regime $\times$ quality interaction term. Within-regime conditional contributions vary substantially: 8.58% within BULL, 28.06% within BEAR, and 13.76% within SIDEWAYS. The BEAR within-regime contribution is the largest of the three, consistent with BEAR's heavier-tailed quality distribution. The substantive contribution is the finding that the analytic decay quantity organises systematically across these coordinates and the coordinates are not redundant. The conditional decomposition shows that quality's contribution does not vanish once regime is conditioned on, but falls from a 30.43% main effect to a 16.80% within-regime contribution. The partial overlap (not an identity) is direct evidence that quality carries information beyond regime rather than being a relabelling of it.

This four-input decomposition is distinct from the three-input total-order decomposition reported in Section 3.2.5 (quality $S_T = 139.7\%$). The three-input total-order index aggregates the within-regime variance and the regime $\times$ quality interaction term under the same column; the four-input decomposition isolates them. Both decompositions are consistent with the same structural reading: quality is regime-conditional. The four-input form supplies the quantitative test; the three-input form supplies the joint-state interaction signature in Section 3.2.5's 307.4% total-order sum.

An 81-cell Black–Scholes Monte Carlo. The 27-cell BS MC of 4.2 extends to 81 cells when regime is added as the fourth axis (3 moneyness $\times$ 3 time $\times$ 3 σ $\times$ 3 regime). Regime-aggregate mean theta values across the 81-cell MC are -40.04% per trading day for BULL (std 42.39%), -41.30% for BEAR (std 42.74%), -39.59% for SIDEWAYS (std 42.16%). The BEAR/SIDEWAYS ratio is $1.04\times$ —directionally consistent with BEAR's elevated-volatility concentration producing marginally faster cliff decay, but the magnitude is modest. The cell-level dispersion within each regime is substantial (std $\approx 42\%$ per regime), reflecting the cliff's joint-state property.

The modest $1.04\times$ BEAR/SIDEWAYS regime-aggregate ratio is worth interpreting with care. Black–Scholes theta on a given (S/K, T, σ) state is invariant to underlying drift; the regime axis enters the 81-cell MC through the sampling distribution over the joint state, not through a theta-modifying parameter. The $1.04\times$ ratio captures only the regime-aggregate shift produced by BEAR's elevated-volatility concentration in the conditional sampling distribution. The substantive regime-conditional structure is in the joint sampling distribution, not in the per-cell theta value.

Joint distribution $P(\text{quality} \mid \text{regime}, \sigma)$. Selected conditional probabilities: SIDEWAYS + Low $\sigma \rightarrow 100.0\%$ LOW quality (deterministic); SIDEWAYS + Medium $\sigma \rightarrow 87.2\%$ LOW; SIDEWAYS + Elevated $\sigma \rightarrow 78.5\%$ MEDIUM, 0.0% HIGH; BEAR + Elevated $\sigma \rightarrow 82.6\%$ HIGH,

0.0% LOW; BEAR + Medium $\sigma \rightarrow 76.4\%$ MEDIUM; BEAR + Low $\sigma \rightarrow 58.6\%$ LOW. The joint distribution is approximately deterministic in three of the (regime, σ) cells documented. Quality cannot be assigned independently of regime and σ —the composite quality score, as calibrated, encodes regime and volatility information through its component factors.

What Section 3.2.7 establishes. Quality is regime-endogenous in the dataset. The regime $\times$ quality contingency-table chi-square ($\chi^2 = 17,823.94$, $p < 0.001$, Cramér's $V = 0.588$) establishes the existence of the dependency at the categorical level. The four-input conditional Sobol decomposition shows a 13.63 percentage-point drop (45% reduction) in quality's variance contribution when regime is included as an explicit input—measuring the size of the dependency at the variance-decomposition level. The joint-distribution assessment confirms the structural endogeneity: in three of the (regime, σ) cells documented, the quality tier is approximately deterministic. The hypothesis is supported. The architectural consequence: framework exit rules should fire on the primary regime signal, not on the derivative quality signal. Future investigations will add a permutation baseline comparison to the conditional Sobol decomposition that would assess the reduction in quality's variance contribution against a random regime label.

What Section 3.2.7 does not establish. The 13.63 pp conditional drop is the variance-retribution signature within the specific 4-input Sobol decomposition deployed here; alternative decompositions (different input axes, different sampling distributions) would produce different point estimates. The 81-cell BS MC reports regime-aggregate theta differences of $1.04\times$ —small in absolute terms, and we do not claim this magnitude is structural. The joint-distribution assessment is based on data in the dataset and conditional on the data window (March–May 2026); whether the same (regime, σ) $\rightarrow$ quality mapping holds in future windows is an open question. The conditional decomposition supplies the formal test of quality endogeneity within the data; it does not test cross-data or cross-time generalisability and does not establish a causal direction (regime $\rightarrow$ quality vs. quality $\rightarrow$ regime).

3.2.8. What the Dataset Establishes and What It Does Not

What the data establishes:

- Sector-conditional cliff timing across 12 sectors with a 7- to 14-day spread between Tech (7–10 DTE, $\rho = +0.849$) and Energy (14–21 DTE, $\rho = -0.958$).
- Quality-stratified decay rates across 150 underlyings with a $6.37\times$ within-cohort HIGH/LOW differential at the effective-theta layer.
- Volatility surface bifurcation by quality with an $8.5\times$ amplitude ratio in the cliff-zone region.
- Variance decomposition, where quality is the largest total-order contributor at $S_T = 139.7\%$.

What the data does not establish:

- A causal claim on quality. The data is observational. The $6.37\times$ differential is consistent with quality-driven mechanisms but does not rule out confounding factors—size, sector mix, idiosyncratic-volatility loadings—that have not been jointly tested in an encompassing regression.
- Subsumption of prior signals. Vasquez's [27] IV-slope conditioning and Cao and Han's [1] idiosyncratic-volatility conditioning may explain some of the same cross-section. Without an encompassing regression, we do not claim quality subsumes either.
- Scope of the three-input total-order indices. The three-input indices (quality $\approx 140\%$ total-order) are reported as an interaction signature, not a precision variance budget, and their three-input first-order companions are not separately tabulated. The first-

order main effect that bears on the regime-endogeneity finding is reported for the four-input decomposition in Section 3.2.7 (Table 2).

- Out-of-sample generalisation. The data covers October 2025 to May 2026 only. Whether the same conditionings hold in subsequent windows is the principal extension question. The tech-positive/energy-negative momentum-correlation signs in Section 3.2.2 reflect the specific regime conditions of the data window: Tech generally flow-positive/energy under war-driven flow-stress. Under different regime conditions—quiet energy markets, technology under correction—the ρ signs may reverse. The $\alpha(s, r)$ parameterisation absorbs this by construction. The specific magnitudes reported here are conditional, not structural.

The conservative reading. Quality conditioning at the effective-theta layer is a measurable feature of the deployment data, complementary to the IV-slope and idiosyncratic-volatility conditionings already in the literature, with subsequent cycles supplying the encompassing-regression test. The $\alpha(s, r)$ parameterisation in three supports a natural extension to $\alpha(s, r, q)$. We do not formalise the extension here. The data evidence is supplementary support for the parameterisation programme, not a separate formal contribution.

4. Discussion

4.1. Implications for Risk Management

The α parameterisation supports an ex ante diagnostic that calendar-based heuristics cannot supply. For each open long-call position, the framework computes the parameter coordinates $(S/K, T, \sigma)$, locates the position on the SIMDEC map, and applies the sector- and regime-conditional $\alpha(s, r)$ to produce an effective theta estimate. Positions in the {NTM, FH, Medium-to-Elevated σ } corner with sector-regime $\alpha > 1$ trigger the rule earlier than calendar-based heuristics would. Positions in other corners with $\alpha < 1$ trigger later. The trigger fires when modelled decay actually crosses the threshold.

Three operational implications follow.

Regional exit rule rather than calendar exit rule. A practitioner who closes positions at fixed days-to-expiry treats the cliff as a clock event. A practitioner who closes positions when their effective-theta coordinates cross into the high-acceleration corner treats the cliff as an input-state event. The latter is closer to the structure of the cliff itself. A position with $T = 5$ days, $S/K = 0.99$, and energy-during-war $\alpha \approx 1.6$ is approaching the corner faster than a calendar rule would suggest. A position with $T = 1$ day, $S/K = 1.10$, and tech-during-quiet $\alpha \approx 0.9$ has substantially less cliff exposure than the calendar would suggest. A calendar rule does not distinguish; the regional rule does.

Book-level α -weighted aggregation. A book with concentrated NTM-FH energy-during-war positions has materially different decay risk than a book with the same total notional spread across the parameter space. The book-level aggregate of α -weighted cliff exposure is observable from the SIMDEC coordinates and the prevailing regime classification. The calendar-based heuristic of “days to expiry weighted by position size” does not capture sector or regime. It misaggregates cliff exposure.

Quantified risk communication. A risk committee receiving a single number—“expected α -weighted cliff loss across the book under the high-acceleration scenario”—can act on a quantified summary. The number is auditable. It follows mechanically from the α table, the SIMDEC coordinates, and the prevailing regime. There is no discretionary judgment between the input data and the reported number.

These are practitioner-relevant uses of the methodology. Whether they translate into measurable risk-adjusted return improvements in a deployed book is an empirical question that the methodology itself does not answer. The methodology supplies the diagnostic. Whether the diagnostic, when applied, produces better outcomes than the

alternative heuristics is a separate question. It requires a longitudinal evaluation comparing α -informed exit rules against calendar-based and price-based alternatives. We identify it as future work.

4.2. Limitations and Future Work

We set out four substantive limitations.

Sample size on the observation. The first observation comes from one cohort on one trading day in one sector under one regime. Three independent positions on a shared parameter-space corner produced indistinguishable decay rates. That is the establishment argument of Section 3.1. What the cohort dataset does not provide is multi-day or multi-cohort coverage. The dataset in Section 3.2 partially addresses this at the population level. It does not substitute for broker-marked replication across multiple expiry dates, multiple sectors, and multiple macroeconomic regimes. The cohort dataset is offered for replication and citation.

Coverage of the parameter space. The cohort sits in one corner of the SIMDEC map. The decomposition's predictions for the other 26 joint-state combinations are not established by the cohort. The dataset extends coverage across 12 sectors and three quality tiers, but does not exhaust the joint-state space. A more complete empirical programme would document positions across each combination and verify that observed decay rates match the regional predictions. We identify this as the most direct extension of the present work.

Single-pipeline data capture. The record was captured by one practitioner via one broker's reporting interface for the cohort, and by one framework's evaluation pipeline for the dataset. The cliff is a property of the option, not of the recording mechanism. Multi-source replication would strengthen the empirical foundation. It should be explicitly recognised that the May 2026 cohort represents a single data point and that the dataset spans a limited window (October 2025–May 2026), so that neither out-of-sample or out-of-period generality can be asserted. Large-scale intraday position-interval testing is considered as part of the continuous-capture work developed in our real-money calibration, companion paper [9]. Irrespective, we invite practitioners with intraday position data from other brokerage environments to contribute analogous trajectories to the public record.

Methodology-to-deployment gap. The implications in Section 4.1 are practitioner-relevant uses of the methodology. They are not measured outcomes from a deployed book. Whether α -informed exit rules produce measurable improvements over calendar-based or price-based alternatives requires a longitudinal evaluation that controls for the comparison rule, the underlying universe, and the operating environment. The companion paper [9] reports a deployment that uses a model-derived edge-decay rule rather than a calendar-based one. The present paper supplies the analytical decomposition that supports such a rule. It does not itself evaluate the deployed-book outcome.

A fifth limitation concerns the rate environment. The simulation assumes $r = 0.045$ as a constant. Sensitivity to large rate moves is not characterised here. Adding r as a fourth input variable to the SIMDEC decomposition expands the joint-state combinations from 27 to 81. The framework supports the extension and the constraint is computational.

A sixth limitation concerns the dataset framing. The data uses framework-computed metrics, not broker marks. The $6.37\times$ decay differential, the $8.5\times$ vol amplification, and the 139.7% Sobol total-order index are data-level findings under the framework's pipeline. Independent replication via a different framework's pipeline, or via direct broker-mark stratification across a comparable universe, would strengthen the data-level claims. We identify this as a research programme requirement.

5. Conclusions

Previous studies have shown sector-conditional asymmetry in implied volatility levels and in option returns. No prior work has parameterised that asymmetry at the effective-theta layer in a form that fires a non-discretionary rule trigger. This paper supplies that parameterisation. Effective theta is defined as $\Theta_e = \alpha(s, r) \cdot \Theta_{BS}$. The SIMDEC filter supplies the methodology for partitioning the input space and identifying the corner where α matters most. The 1 May 2026 cohort supplies a first observation: three concurrent long-call positions in the predicted corner, decaying at indistinguishable rates within 1 percentage point of each other across the most extreme interval, with cohort-mean $\hat{\alpha} \approx 1.6$ for energy-during-war.

The SIMDEC dataset from the same deployment supplies population-level support across 12 sectors and three quality tiers, with 25,789 records across 8 states carrying the framework's BULL/BEAR/SIDEWAYS regime classification. Sector-conditional cliff timing spans 7- to 14-day windows, with Tech at 7–10 DTE and Energy at 14–21 DTE. Quality stratification produces a $6.37\times$ HIGH/LOW decay-rate differential at the effective-theta layer and an $8.5\times$ amplitude ratio at the volatility surface. Variance decomposition under total-order Sobol assigns the largest single contribution to quality. A conditional decomposition with regime as an explicit fourth input shows a 13.63 percentage-point drop in quality's variance contribution, confirming quality's regime endogeneity. The dataset evidence is complementary to the 1 May cohort. It does not substitute for it.

We position the May 2026 observation as a first illustration and frame the dataset as supplementary illustrative evidence at the population level. This represents a single-cycle illustration, not a generalisable performance claim. The token-free analytical formulation and the corresponding explainability methodology within regulated decision-making settings represent the durable contribution. The first observation shows that the formulation has at least one cycle of evidence in its predicted corner. The dataset shows the formulation has cross-sectional support consistent with sector- and quality-conditional structure. Subsequent cycles establish the magnitude.

The dataset is offered for replication, citation, and further investigation. The cliff is the cliff regardless of who captures the data. The cliff is a regional phenomenon in the parameter space, regardless of which session it is observed in. The most significant business contributions from the framework are the α parameterisation at the effective-theta layer and the methodology for calibrating it. Both of these contributions are useful independently of the framework architecture used to capture the record.

6. Patents

The continuous-snapshot data-capture methodology referenced in Section 3.1.2 is one component of a broader portfolio-management framework. That framework is the subject of a pending patent application filed by the author, G.M., through Innovation York at York University. Author G.M. is the named inventor. The data observations, the SIMDEC decomposition, and the α parameterisation reported in this paper do not depend on the framework's other components and are entirely reproducible, in full, from any equivalent broker-CSV-export workflow.

Author Contributions: G.M. and J.Y. were involved throughout the entire process of producing this paper, including conceptualisation, methodology, modelling, validation, visualisation, and manuscript preparation. All authors have read and agreed to the published version of the manuscript.

Funding: This research was supported in part by grant OGP0155871 from the Natural Sciences and Engineering Research Council. This study received no external corporate funding or third-party

capital. The data-capture deployment that generated the empirical record used the author G.M.'s personal capital held in regulated retail brokerage accounts.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The eight-snapshot trajectory dataset described in Section 3 is held by the author, G.M., and is available to qualified academic researchers under reviewer-confidentiality protocols. The dataset includes timestamps, broker-recorded marked prices, position quantities, and cohort-level attribution for the three positions across the eight distinct intraday states. Public-record references for the macroeconomic context are independently verifiable and are cited in Appendix A.

Acknowledgments: During the preparation of this manuscript, the authors used Claude (Anthropic) for language polishing, structural editing, and literature search support. The authors have reviewed and edited the output and take full responsibility for the content of this publication. The authors would like to thank the Editor and three anonymous reviewers for providing numerous helpful suggestions that have substantially improved the overall content and delivery of the paper.

Conflicts of Interest: The authors declare no conflicts of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Abbreviations

The following abbreviations are used in this manuscript:

AI/ML	Artificial Intelligence/Machine Learning
AMD	Advanced Micro Devices Inc., NASDAQ
ATM	At-the-money
BiLSTM	Bidirectional Long Short-Term Memory
BS	Black–Scholes
CAD	Canadian Dollar
CVX	Chevron Corp, NYSE
DTE	Days-to-expiration
FH	Final Hours
FS	Final Session
ITM	In-the-money
LNG	Chevron Corp, NYSE
MC	Monte Carlo
MS	Multi-Session
NTM	Near-the-money
OTM	Out-of-the-money
QMJ	Quality Minus Junk
USD	United States Dollar
XOM	Exxon Mobil Corp, NYSE
XAI	Explainable Artificial Intelligence

Appendix A

Macroeconomic Context of the 1 May 2026 Session

The illustrative dataset was captured during the cliff session of 1 May 2026. The session fell within an active geopolitical and energy-market volatility regime. The macroeconomic context appears here for completeness. It is not a load-bearing element of the methodological contribution. The cliff is a regional phenomenon in the SIMDEC parameter space. It would have presented identically on a different expiry date in a quieter regime, provided the parameter coordinates were the same.

The 2026 Iran war began on 28 February 2026 with coordinated United States and Israeli airstrikes [28,29]. On 4 March 2026, Iran declared the Strait of Hormuz closed to vessels of the United States, Israel, and Western allies. The International Energy Agency characterised the resulting disruption as a substantial supply disruption in the global oil market. A ceasefire announced on 8 April 2026 was followed by a renewed Hormuz closure on or around 26 April 2026. Brent crude prices traded across approximately USD 80–82 per barrel in early March 2026 to approximately USD 112 per barrel by late April 2026. The first U.S. airline casualty of the sustained fuel-price shock was Spirit Airlines. The shutdown was announced on 2 May 2026—one day after the cliff session reported here [30,31].

The implication for the data record is that the cohort's three underlyings (XOM, CVX, LNG) traded in a moderately elevated implied-volatility regime during the validation window. This places σ in the Medium-to-Elevated bin of the SIMDEC decomposition. The placement is consistent with the high-acceleration corner identified in Section 2.4.3. The observed $\hat{\alpha} > 1$ matches the expected sign for energy during a stress regime as defined in Section 2.3. A quieter volatility regime would place the cohort in a different σ bin. The decomposition would then predict a less extreme cliff. The absolute decay would be smaller. The relative cliff structure would be the same.

Appendix B

Appendix B.1. α Back-Out: Computation from the Eight-Snapshot Trajectory

This appendix supplies the per-position, per-interval α back-out from the eight-snapshot cohort trajectory of Section 3.1. The computation follows the Section 2.4.3 formula. Observed effective theta is computed from the broker-marked decay across each snapshot interval. Black–Scholes theta is computed from the position parameters at the same timestamps using the closed-form expression of Section 2.2.

Appendix B.1.1. Per-Position Computation Method

For each position and each snapshot interval $[t_i, t_{i+1}]$, observed effective theta is

$$\hat{\Theta}_e^{obs} = \frac{V(t_{i+1}) - V(t_i)}{t_{i+1} - t_i}$$

where $V(t)$ is the broker-marked value at timestamp t . Black–Scholes theta is computed from the position parameters $(S(t_i), K, T(t_i), \sigma(t_i), r)$ at the start of the interval. The interval α estimate is

$$\hat{\alpha}_i = \frac{\hat{\Theta}_e^{obs}}{\Theta_{BS}(t_i)}.$$

For the first cliff acceleration interval, T1–T2 (approximately 09:35 EDT to 10:55 EDT), the per-position estimates appear in Table A1.

Table A1. Per-position α back-out across the T1–T2 interval, 1 May 2026.

Position	T1 Mark (USD)	T2 Mark (USD)	Observed Decay (% Premium)	BS Theta-Implied Decay (% Premium, Approximate)	$\hat{\alpha}_{T1-T2}$ (Approximate)
XOM 160C	0.45	0.07	84.4%	50–55%	~1.6
CVX 175C	0.85	0.14	83.5%	48–52%	~1.7
LNG 250C	0.70	0.11	84.3%	45–50%	~1.7

The Black–Scholes-implied decay rates over the T1–T2 interval are computed at the corresponding underlying spots, the position strikes, the time-to-expiry implied by the snapshot timestamps, and the contemporaneous implied volatility on each po-

sition. The reported ranges reflect uncertainty in the precise IV input at each timestamp. Broker-displayed mid-mark IV and the implied IV from the marked premium can differ at the second-decimal level. The reported $\hat{\alpha}$ values are interval estimates within the reported ranges.

The cohort-mean estimate for the T1–T2 interval is approximately $\hat{\alpha} \approx 1.6$ to 1.7. The interpretation: during the war-regime energy-sector environment of 1 May 2026, observed decay on the cohort exceeded Black–Scholes theta predictions by approximately 60–70% across the most extreme interval of the cliff trajectory.

Appendix B.1.2. Subsequent-Interval Behaviour

The per-interval $\hat{\alpha}$ estimates fall toward unity in the latter intervals of the session. By the T8–T9 interval, the cohort positions are at terminal premium values where the broker-marked decay rate is bounded below by the rounding precision of broker quotes (USD 0.01). The $\hat{\alpha}$ estimate becomes mechanical rather than informative in those intervals. The methodologically informative intervals are T1–T2 and T2–T3, where premium remains above the rounding floor and broker marks track underlying movement at the mid-mark cadence.

Appendix B.1.3. Reproducibility Notes

The overall method is highly reproducible. SimDec is completely open-source on GitHub, allowing any business to run it locally and deterministically, bypassing the need for proprietary data centres (and from incurring astronomical cloud-based token costs). All SIMDEC code is freely available, open-source, in Python, R, Matlab, and Julia from <https://github.com/Simulation-Decomposition> (accessed on 9 July 2026). A no-code-required, web-based dashboard can be accessed at <https://simdec.io/> (accessed on 9 July 2026). An open access SIMDEC reference book can be downloaded from: <https://doi.org/10.4324/9781003453789>. The authors can supply replication code in Python or R to researchers on request. The code is independent of the data-capture pipeline referenced in the Conflicts of Interest section. It uses only the closed-form Black–Scholes pricer, standard implied-volatility solvers, and the broker-CSV record reported in Table 1. The cohort-level dataset is held by the authors. It is available to qualified academic researchers under reviewer-confidentiality protocols. Researchers, practitioners, and companies can readily test their own specific trading strategies (whether proprietary or open-source) by straightforwardly incorporating them directly within the architecture of our prescribed framework. Hence, the methodology is readily reproducible.

The central claim of groundedness-by-construction does not depend on access to any restricted element. The property follows from the disclosed architecture, which contains no generative step from which an out-of-cache output could arise and can be straightforwardly verified against output logs. There is a distinction that needs to be made explicit between released code, re-runnable analysis, and a re-implementable specification. The currently released code is the SimDec package, which is completely open-source and publicly available on GitHub. The SimDec filtering is fully re-runnable because the data layer draws exclusively on free public APIs—no access to the broker record and no proprietary code is required. The architectural pattern is specified in the paper and can be independently re-implemented with any MIP solver (including open-source packages). The data layer draws exclusively on free public APIs (Yahoo Finance market data, Bloomberg sector taxonomy, EIA petroleum statistics) and these together make the SimDec filtering fully re-runnable by independent researchers with no access to the broker records. The architectural pattern is fully disclosed at a level sufficient for independent re-implementation with any MIP solver. Primary-source record excerpts (framework dashboards, trade tables, timestamps,

and broker portfolio summary figures) are reproduced in the manuscript (or in [9]) and constitute the paper's documentary evidence base. The specific investment outcomes are not replicable, but no reader is asked to replicate them. These are operator-, capital-, and period-specific, as stated, and no claim on the methodological approach depends upon their replication.

A scoping note. The $\hat{\alpha}$ values reported here are interval estimates from one cohort on one trading day. They are a first observation. They are not magnitude estimates for $\alpha(s = \text{energy}, r = \text{war})$ across the population of trading sessions in that sector-regime combination. The α parameterisation in Section 3 supplies the framework for calibrating the population estimate as future cycles supply observations. The 1 May 2026 cohort supplies one observation. The mathematical formulation does not depend on the magnitude estimate. The first observation shows the formulation has at least one cycle of evidence in its predicted corner.

References

1. Cao, J.; Han, B. Cross section of option returns and idiosyncratic stock volatility. *J. Financ. Econ.* **2013**, *108*, 231–249. [CrossRef]
2. Natenberg, S. *Option Volatility and Pricing: Advanced Trading Strategies and Techniques*, 2nd ed.; McGraw-Hill: Columbus, OH, USA, 2015.
3. Hull, J.C. *Options, Futures, and Other Derivatives*, 10th ed.; Pearson: London, UK, 2018.
4. Black, F.; Scholes, M. The pricing of options and corporate liabilities. *J. Political Econ.* **1973**, *81*, 637–654. [CrossRef] [PubMed]
5. Carr, P.; Wu, L. Variance risk premiums. *Rev. Financ. Stud.* **2009**, *22*, 1311–1341. [CrossRef]
6. Bakshi, G.; Kapadia, N. Delta-hedged gains and the negative market volatility risk premium. *Rev. Financ. Stud.* **2003**, *16*, 527–566. [CrossRef]
7. Christoffersen, P.; Heston, S.; Jacobs, K. The shape and term structure of the index option smirk: Why multifactor stochastic volatility models work so well. *Manag. Sci.* **2009**, *55*, 1914–1932. [CrossRef]
8. Kozlova, M.; Yeomans, J.S. *Sensitivity Analysis for Business, Technology, and Policymaking: Made Easy with Simulation Decomposition (SimDec)*; Routledge: Abingdon-on-Thames, UK, 2024. [CrossRef]
9. Melville, G.; Yeomans, J.S. Filter Before You Solve: A Deterministic-First/Learned-Second Architecture for AI-Driven Portfolio Management with Real-Money Training-Investment Calibration. In *Working Paper*; National Bureau of Economic Research: Cambridge, MA, USA, 2026. [CrossRef]
10. Kozlova, M.; Ahola, A.; Roy, P.; Yeomans, J.S. Simple Binning Algorithm and SimDec Visualization for Comprehensive Sensitivity Analysis of Complex Computational Models. *J. Environ. Inform. Lett.* **2025**, *13*, 38–56. [CrossRef]
11. Schwab Center for Financial Research. Theta Decay in Options Trading: Three Strategies. Charles Schwab. 2026. Available online: <https://www.schwab.com/learn/story/theta-decay-options-trading> (accessed on 9 July 2026).
12. CME Group. Theta and Option Pricing. CME Group Education Materials. 2024. Available online: <https://www.cmegroup.com/education/courses/option-greeks/theta> (accessed on 9 July 2026).
13. Bakshi, G.; Cao, C.; Chen, Z. Empirical performance of alternative option pricing models. *J. Financ.* **1997**, *52*, 2003–2049. [CrossRef]
14. Andersen, T.G.; Bollerslev, T.; Diebold, F.X.; Labys, P. Modeling and forecasting realized volatility. *Econometrica* **2003**, *71*, 579–625. [CrossRef]
15. Barndorff-Nielsen, O.E.; Shephard, N. Econometric analysis of realized volatility and its use in estimating stochastic volatility models. *J. R. Stat. Soc. Ser. B* **2002**, *64*, 253–280. [CrossRef]
16. Ni, S.X.; Pearson, N.D.; Poteshman, A.M. Stock price clustering on option expiration dates. *J. Financ. Econ.* **2005**, *78*, 49–87. [CrossRef]
17. Kozlova, M.; Moss, R.J.; Yeomans, J.S.; Caers, J. Uncovering heterogeneous effects in computational models for sustainable decision-making. *Environ. Model. Softw.* **2024**, *171*, 105898. [CrossRef]
18. Tarantola, S.; Ferretti, F.; Piano, S.L.; Kozlova, M.; Lachi, A.; Rosati, R.; Puy, A.; Roy, P.; Vannucci, G.; Kuc-Czarnecka, M.; et al. An Annotated Timeline of Sensitivity Analysis. *Environ. Model. Softw.* **2024**, *174*, 105977. [CrossRef]
19. Sobol, I.M. Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates. *Math. Comput. Simul.* **2001**, *55*, 271–280. [CrossRef]
20. Saltelli, A.; Ratto, M.; Andres, T.; Campolongo, F.; Cariboni, J.; Gatelli, D.; Saisana, M.; Tarantola, S. *Global Sensitivity Analysis: The Primer*; John Wiley & Sons: Hoboken, NJ, USA, 2008.
21. Helo, J.; Kozlova, M.; Roy, P.; Yeomans, J.S. Regional Variance-Based Sensitivity Analysis for Complex Models. In *Working Paper*; National Bureau of Economic Research: Cambridge, MA, USA, 2026.

22. Bates, D.S. Jumps and stochastic volatility: Exchange rate processes implicit in Deutsche Mark options. *Rev. Financ. Stud.* **1996**, *9*, 69–107. [\[CrossRef\]](#)
23. Figlewski, S. What does an option pricing model tell us about option prices? *Financ. Anal. J.* **1989**, *45*, 12–15.
24. Heston, S.L.; Jones, C.S.; Khorram, M.; Li, S.; Mo, H. Option momentum. *J. Financ.* **2023**, *78*, 3141–3192. [\[CrossRef\]](#)
25. Asness, C.S.; Frazzini, A.; Israel, R.; Moskowitz, T.J.; Pedersen, L.H. Size matters, if you control your junk. *J. Financ. Econ.* **2018**, *129*, 479–509. [\[CrossRef\]](#)
26. Asness, C.S.; Frazzini, A.; Pedersen, L.H. Quality minus junk. *Rev. Account. Stud.* **2019**, *24*, 34–112. [\[CrossRef\]](#)
27. Vasquez, A. Equity volatility term structures and the cross section of option returns. *J. Financ. Quant. Anal.* **2017**, *52*, 2727–2754. [\[CrossRef\]](#)
28. CSIS. *Operation Epic Fury and the Remnants of Iran's Nuclear Program*; Center for Strategic and International Studies (CSIS): Washington, DC, USA, 2026. Available online: <https://www.csis.org/analysis/operation-epic-fury-and-remnants-irans-nuclear-program> (accessed on 9 March 2026).
29. U.S. Department of War. Operation Epic Fury: First 10 Days. 2026. Available online: <https://media.defense.gov/2026/Mar/09/2003896756/-1/-1/1/OPERATION-EPIC-FURY-FACT-SHEET-THE-FIRST-10-DAYS.PDF> (accessed on 9 March 2026).
30. CNN Business. Spirit Airlines Canceled All Flights and Is Going Out of Business. 2026. Available online: <https://www.cnn.com/2026/05/02/business/spirit-to-halt-all-flights> (accessed on 2 May 2026).
31. Reuters. *Spirit Airlines Shuts Down, Industry's First Iran War Casualty*; Thomson Reuters: New York, NY, USA, 2026.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.