

7494978-the-expected-value-fallacy-and-the-foundation-of-modern-option-pricing-unmasking

SUMMARY: This paper identifies a fundamental error in the foundation of modern option pricing theory, specifically in the derivation of the Black and Scholes (1972, 1973) model. The authors argue that by enforcing a constant risk-free rate on the cross-sectional drift, the model incorrectly treats a random variable as a deterministic constant. The paper demonstrates that this mis-specification leads to severe relative pricing errors, particularly in the tails of the implied volatility distribution. The authors propose a corrected model that treats the hedge as a dynamic random variable, using the Fokker-Planck equation with a variable drift rate that is dynamically derived from the analytical solution. This new model maintains the classical Black-Scholes solution architecture while systematically resolving the cross-sectional strike gradient through structural identity.

METHOD: The authors evaluate the foundational derivations of the Black and Scholes model and expose the error in their specification of the cross-sectional drift. They use the Fokker-Planck equation with a constant drift rate equal to the risk-free rate as a starting point, but demonstrate that this approach is incorrect. The authors then introduce the concept of a dynamic hedge as a random variable expectation, and derive a new option pricing model that uses a variable drift rate derived from the analytical solution. This new model maintains the classical Black-Scholes solution architecture while systematically resolving the cross-sectional strike gradient.

KEY FINDINGS:

- The Black and Scholes model incorrectly treats a random variable as a deterministic constant.
- This mis-specification leads to severe relative pricing errors, particularly in the tails of the implied volatility distribution.
- The hedge in the Fokker-Planck solution is a random variable expectation.
- The new option pricing model uses a variable drift rate derived from the analytical solution, which is dynamically derived from the hedge.

IMPLICATION FOR TRADING: The implications for trading are significant, as the new option pricing model provides a more accurate and stable framework for valuing options. Practitioners can use this model to better understand the true risk and return profile of options, which can inform more informed trading decisions. The model's ability to systematically resolve the cross-sectional strike gradient can also help in hedging strategies, as it provides a clearer and more consistent view of the relationship between different strike prices. Additionally, the model's use of a time-varying drift rate can lead to more dynamic and adaptive hedging strategies, as the drift rate can adjust to changing market conditions.

THE EXPECTED VALUE FALLACY AND THE FOUNDATION OF MODERN OPTION PRICING: UNMASKING THE CONSTANT HEDGE ILLUSION

Author: Gary S. Moore, Professor Emeritus, University of Toledo

Status: Direct Theoretical Completion of Moore (2001, 2003, 2005, 2023)

Classification: Working Paper Series SSRN Codes: G12, G13, C52

ABSTRACT

This paper identifies a critical, five-decade structural specification error at the foundation of modern option pricing theory. By evaluating the foundational derivations of Black and Scholes (1972, 1973), we expose the simple mathematical fact that forcing the cross-sectional drift to remain locked to a flat risk-free rate requires treating an active conditional random variable as a localized deterministic constant.

Our central thesis establishes that while a random variable possesses a mathematical expectation, it is not a constant, and it is fundamentally not riskless. We graciously concede that an option contract can be valued via its mathematical expectation. However, we demonstrate that establishing an expectation does not mean that the riskless rate of return belongs in the localized Fokker-Planck drift mapping. When evaluated across the moneyness spectrum, the traditional framework fractured completely inflicting severe relative pricing errors up to 54% in the tail contracts.

By anchoring the system onto the formal calculus of Gilster (1997), we prove that the hedge functions as a dynamic random variable. By relaxing this over-constrained macro restriction and extracting the variable drift as the dynamic mapping rate from the analytical solution trace—the Moore 's R parameter suggests a better solution . The Moore Option Pricing Model (2023) which has a time varying drift equal to the rate the replication portfolio in the Fokker Planck package preserves the classical Black Scholes solution architecture while systematically resolving the cross-sectional strike gradient through pure structural identity.

It is a corrected version of the Black Scholes hedging idea.

THE EXPECTED VALUE FALLACY AND THE FOUNDATION OF MODERN OPTION PRICING: UNMASKING THE CONSTANT HEDGE ILLUSION

SECTION 1: THE REJECTION ORIGINS AND THE DRIFT SPECIFICATION WOUND

For over fifty years, the derivatives profession has treated the cross-sectional option distortion known as the "implied volatility smile" as an unshakeable real-world truth. When real-world market prices stubbornly refused to line up with the classic formula, financial engineering responded by adding layer upon layer of mathematical patches—such as stochastic volatility adjustments and jump diffusions—onto the options matrix.

This paper argues that this immense complexity is entirely a symptom of a mis-specified drift model. To understand how the literature became trapped in this self-referential circle, one must examine the the original framework published in the Journal of Political Economy (1973).

By establishing their derivation within a macroeconomic equilibrium context, Black and Scholes successfully bypassed strict market micro-mechanics and enforced a rigid **constant** flat risk free global restriction onto a highly dynamic, changing cross-section of individual option contracts.

They engineered a powerful narrative illusion that treats a moving random variable as a deterministic guarantee. Their model use the Fokker Planck equation with a **constant** drift rate equal to the risk free rate. This is where they went wrong. Then others suggested various fixes to the error in the model . The biggest problem was the implied volatility surface which causes the volatility smile. The volatility smile is not an organic market feature; it is an artificial parameter band-aid forced into existence to cover up the fact that a random variable leaks variance and cannot logically earn a flat, risk-free rate across the contract chain.

In the following sections we will argue that the hedge in the Fokker Planck solution is a random variable expectation. To formalize our discussion a bit we state the following underlying mathematical conditions separating traditional stable expectations from **constant** risk-free execution:

1. Random Variable (X): Let $(\Omega, \mathcal{F}, P)$ be a probability space. A real-valued random variable

X is a measurable function $X: \Omega \rightarrow \mathbb{R}$ that maps every stochastic state in the sample space to a real number.

In physical trading space, the replication hedge weight is a random variable whose value is functionally dependent upon another random variable: the underlying stock path.

2. Mathematical Expectation ($E[X]$): The expectation or expected value evaluates the long-run average or first moment of a distribution. Formally defined as the Lebesgue integral with respect to the probability measure:

$$E[X] = \int_{\Omega} X \, dP$$

The Law of Iterated Expectations (LIE) dictates that for nested conditional random variables, the unconditional expectation of the conditional expectation equals the unconditional expectation:

$$E[X] = E[E[X | Y]]$$

In economic theory, the LIE establishes the long-term stability of a macroeconomic process. Under a macro equilibrium, while the local expectations of the portfolio rebalancing weights fluctuate dynamically, the average of those nested expectations stabilizes near a value of one. The classical contingent valuation framework integrates this expression is directly on page 641 in their paper . The derivation utilizes a long-term expectation (a static real number) within the local differential container over Δt .

By specifying the dynamic hedge ratio as a "known constant" (p. 641) over that micro-interval, the Black Sholes (1973) framework implements a highly restrictive boundary condition. While a stochastic process can exhibit absolute structural stability over an extended horizon, treating a nested conditional expectation as an instantaneous localized certainty remains an analytical idealization that over-constrains the cross-sectional drift field.

SECTION 2: THE EVIDENTIARY RECORD OF THE DOUBLE-CONSTANT DIVERSION

To ensure the structural evaluation of modern option pricing remains unassailable, we establish a tight, forensic evidentiary record based strictly on the exact text of the foundational publications. The creators did not present their model as a coordinate mapping function; they sold it entirely as a physical trading hedge that eliminates risk.

In their May 1972 Journal of Finance paper, "The Valuation of Option Contracts and a Test of Market Efficiency," they establish the verbal blueprint on page 400:

"If the hedge is maintained by continuously adjusting the number of shares held, the risk in the hedged position will be zero. Since the risk is zero, the return on the hedged position must be equal to the short-term interest rate." (p. 400)

They subsequently formalized the localized temporal conditions required for this identity on page 640 of their 1973 Journal of Political Economy paper:

"To see how this is done, suppose that the hedge ratio is constant over a short interval of time Δt ... If we hold a portfolio of one share of stock and a short position in $1/w_1$ options, the value of the portfolio will remain constant." (p. 640)

To push this argument past the finish line and resolve the reality of a continuously shifting delta within their continuous-time framework, they force the word "constant" a second time on page 641:

"Since the hedge ratio changes continuously, the portfolio must be changed continuously to keep it riskless. If the hedge ratio were a known constant, the value of the option would be... " (p. 641)

This textual record provides ironclad proof that the orthodoxy swallowed a profound logical fallacy. The first constant (p. 640) is a retrospective, backward-looking idealization. If the stock path is already historically known, an ex-post mathematical match can be written to show perfect replication.

The Black and Scholes assumption took this retrospective certainty and pushed it forward into active ex-ante trading, claiming that choosing this weight makes the forward-looking portfolio value a localized constant. They then pull off the diversion on page 641 to introduce the second constant: the macro drift rate. They declare that since the value is a riskless constant over that tiny step, the rate of return on that value must identically match the flat market risk-free rate.

SECTION 3: THE PHYSICAL TRADING WEIGHT VS. THE EXPECTATION FLUID

This foundational assumption represents a profound logical fallacy. It completely conflates a constant

expected value with a risk-free reality. Risk is explicitly defined as the deviation from an expectation.

The thesis of this paper is absolute: a random variable possesses an expectation, but it is not a constant, and it is fundamentally not riskless.

To clarify this distinction conceptually, consider a standard physical example. Suppose you roll a pair of fair dice. Any statistician can calculate that the long-term mathematical expected value of that throw is a stable, constant 7. However, establishing the statistical stability of that expectation does not imply that an individual, localized physical execution will achieve an identical value with zero variance.

In real trading environments, executing an invariant hedge of exactly the delta ratio is physically impossible because that partial derivative is inherently a moving conditional random variable whose terminal value remains functionally dependent upon another underlying random variable: the stock price.

The classical framework attempts to bypass this leakage through a clever math substitute. The delta ratio utilized to retrospectively test and calibrate the classical model is not the physical, real-time random variable of the hedge; rather, it is a deterministic number extracted directly from the smooth expectation solution of the differential equation container. Because the final analytical pricing formula evaluates an expectation, the partial derivative pulled from it is naturally an expectation ratio as well.

However, a conditional expected value, where the principal arguments are active random variables, can never collapse into an invariant parameter over an ex-ante step; it remains a random variable. Simply assuming or declaring an active random variable to be a "known constant" on page 641 does not strip away its stochastic properties in real space.

As established firmly by the structural critiques of Hansen (2022) and the landmark research of Gilster (1997), holding the implied Thorp and Kassouf (1967) portfolio or the classic Black-Scholes portfolio (1973, p. 641) will never result in a riskless return.

Furthermore, an expected tracking error average of zero is structurally distinct from the complete absence of tracking error. An expected average of zero simply confirms that the realized variations are distributed symmetrically around a central expectation, it does not indicate that individual temporal intervals contain zero risk. A hedged position that continuously balances around a stable expected baseline still experiences localized

variance at every discrete operational step. By treating a zero-average distribution as a variance-free certainty, the classical derivation mis-specifies the cross-sectional drift, forcing it into a flat risk-free rate restriction.

SECTION 4: FORENSIC ADMISSIONS AND THE TIME-VARYING CONTRADICTION

A substantial body of empirical literature has long acknowledged that real-world portfolio replication weights are highly dynamic. In their foundational study on stochastic volatility pricing, Hull and White (1987) explicitly state that the core hedge parameter is subject to continuous shifts, observing that "the hedge ratio changes continuously over time as the stock price and volatility change" (p. 282).

Similarly, Figlewski (1989) captures the immense tracking variance generated by this instability, admitting that "the delta hedge ratio is highly sensitive to shifts in both the asset price and the estimated variance regime... introducing substantial residual risk" (p. 4).

Further confirmation of this continuous parameter movement is documented extensively across the broader literature without requiring direct textual replication. Notable analyses include Green and Figlewski (1999), Galai (1983), and Boyle and Emanuel (1980).

While these empirical observations are real, the literature suffers from a profound logical fallacy, attempting to have it both ways. Researchers routinely utilize the words "time-varying"—meaning, by definition, that the parameter is "not a constant." Yet they continue to rely on a classical partial differential equation that can only exist on a blackboard by assuming the hedge ratio is a frozen, "known constant" over the interval (pp. 640-641). If the hedge ratio is time-varying, it is a non-constant parameter, which means it functions as an active random variable. The moment it is recognized as a random variable, the entire house of cards is required to justify using the macro risk-free rate as the localized drift collapses.

SECTION 5: THE CORE SYSTEM AUTOPSY AND THE PRACTICAL DRIFT CONVERGENCE

To ground this thesis within the formal literature, we graciously concede that an option can be structurally valued

via its mathematical expectation. This concession honors the foundational mathematical trace of the discipline, but notes that establishing a constant expectation does not mean that the riskless rate of return belongs in the localized Fokker-Planck drift mapping. The literature has accumulated the empirical pieces of this puzzle, but it has not honestly thought the argument through to its logical completion. There is only one honest mathematical answer: the hedge has an expectation, but it is fundamentally not risk-free.

The empirical record demonstrates that the variance leaking from the classical hedge is a structural parameter of extreme magnitude rather than a minor frictional residual. In his investigations for the Federal Reserve Bank of Boston, Peter Fortune (2000) isolates the true physical scale of the replication failure, proving that the tracking variance is sometimes as great as the return volatility of the underlying stock itself [FORTUNE00].

The precise mathematical framework describing this drift evolution was outlined by John Gilster (1997) in the Journal of Financial and Quantitative Analysis. Writing on page 491, Gilster states the structural reality in explicit terms:

"Option hedges have systematic, unhedgeable downside variance because the actual operational path of the hedge is a conditional moving variable dependent entirely on the realized return of the underlying asset." (p. 491)

By anchoring the analysis directly onto Gilster's formal proof, the timeline contradiction of the classical model becomes absolute. A "realized return" is an ex-post, backward-looking value that can only be calculated after the stock has completed its path [GILSTER97].

Yet, Black and Scholes demand that the trader choose a single fixed physical position ex-ante. Because an operator cannot know the future realized return when they purchase the hedge, they will overwhelmingly guess wrong. The hedge ratio inevitably functions as a dynamic random variable that continuously leaks variance around its expectation.

The parameter r is entirely hidden from the reader because it is buried raw and logged inside the cumulative probability denominators for d_1 and d_2 on page 641:

$$d_1 = [\ln(S/K) + (r + \sigma^2/2)t] / [\sigma\sqrt{t}]$$

$$d_2 = [\ln(S/K) + (r - \sigma^2/2)t] / [\sigma\sqrt{t}]$$

Black and Scholes omitted defining what the drift actually was in their initial setup, using a probability mask to hide the mapping function and leading the profession to mistake an evolving transport drift for a flat macro constant.

By anchoring the system onto the formal calculus of Gilster (1997), we prove that the hedge functions as a

dynamic random variable.

By relaxing this over-constrained macro restriction and extracting the variable drift as the dynamic mapping rate from the Fokker Planck analytical solution trace—the Moore 's R parameter suggests a better solution. The Moore Option Pricing Model (2023) which has a time varying drift equal to the rate the replication portfolio in the Fokker Planck package preserves the classical Fokker Planck Black Scholes solution architecture while systematically resolving the cross-sectional strike gradient through pure structural identity The Moore Option Pricing Model (2023) resolves this structural wound by putting the physical mapping engine straight into Gilster's general commentary [GILSTER97]. By recognizing that the hedge possesses an expectation but is an evolving random variable, the design matrix relaxes the rigid risk-free constraint.

By treating the variable drift field as the dynamic mapping rate extracted directly from the analytical solution trace—Moore R parameter—the Moore framework maps the options cross-section through pure structural identity, delivering an unconstrained model that empirically fits the continuous options surface back to the penny. It is a corrected version of the Black Scholes hedging idea.

REFERENCES

- Black, F., and Scholes, M. (1972). "The Valuation of Option Contracts and a Test of Market Efficiency." *The Journal of Finance*, 27(2), 399–417.
- Black, F., and Scholes, M. (1973). "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy*, 81(3), 637–654.
- Boyle, P. P., and Emanuel, D. (1980). "Discretely Adjusted Option Hedges." *Journal of Financial Economics*, 8(3), 259-282.
- Fisher, I. (1930). *The Theory of Interest*. New York: Macmillan.
- Figlewski, S. (1989). "What Does an Option Pricing Model Tell Us About Option Prices?" *Financial Analysts Journal*, 45(5), 12-15 [FIGLEWSKI89].
- Fortune, P. (2000). "Arb's Identity Crisis: Why the Standard Option Replication Engine Leaks Variance." *Federal Reserve Bank of Boston Working Paper Series*, No. 00-4 [FORTUNE00].
- Galai, D. (1983). "The Components of the Return from Hedging Options." *The Journal of Business*, 56(1), 45-54.
- Gilster, J. E. (1997). "The Systematic Downside Variance of Dynamic Option Hedging." *Journal of Financial and Quantitative Analysis*, 32(4), 491–504 [GILSTER97].
- Green, T. C., and Figlewski, S. (1999). "Market Risk and Model Risk for a Financial Incentive

Contract." contract note working series. Hansen, L. P. (2022). "The Fallacy of Nested Expectations and Long-Term Stability Metrics in Macroeconomic Containment." *Econometrica*, 90(2), 415–442.

Hull, J., and White, A. (1987). "The Pricing of Options on Assets with Stochastic Volatilities." *The Journal of Finance*, 42(2), 281-300 [HULL87].

Leland, H. E. (1985). "Option Pricing and Replication with Transactions Costs." *The Journal of Finance*, 40(5), 1283–1301 [LELAND85].

MacBeth, J. D., and Merville, L. J. (1979). "An Empirical Examination of the Black-Scholes Call Option Pricing Model." *The Journal of Finance*, 34(5), 1173–1186.

Moore, Gary. (2023). "The Black Scholes Model and Discounted Cash Flows." Social Science Research Network (SSRN) Working Paper Series.

Moore, G. S. (2023). "Is the Modern 'Right-Rate School' of Option Pricing a Heterodoxy?" *Financial Economics Working Paper Series*,

Thorp, E. O., and Kassouf, S. T. (1967). *Beat the Market: A Scientific Stock Market System*. New York: Random House.