RESEARCH MEMO

The Covered Call Overlay

What you are actually selling, where the premium comes from, and why the strategy is judged on the wrong axis

Steven Ede, CFA

September 2026

Modeled framework. Analytic results are derived in full; distributional results are specified in §10 and marked where outputs are pending. No proprietary data is used.

1 · SUMMARY OF CONCLUSIONS

A covered call program is usually presented as an income strategy and evaluated on total return. The first description is incomplete rather than wrong; the second is a genuine error.

The income is real. Premium is cash, received on a schedule, and funding distributions from it rather than by selling shares is a legitimate and useful application. But income is the form the return takes, not its source. The source is a short volatility position harvesting a specific and well-documented risk premium, and a program described only by its cash flow tends to get evaluated on total return in a rising market, a comparison it must lose by construction, because capping upside is what it does.

Judged on the axes that matter, return per unit of downside risk, reliability of distributions, and whether the client has any use for the extreme right tail, the overlay earns a place in more portfolios than it currently occupies.

> Treat the overlay as a harvest of the variance and correlation risk premium, sized against a specific underlying the client intends to hold regardless, and evaluate it against a downside-sensitive benchmark rather than total return.

Findings

1 · The excess return comes entirely from a volatility wedge. Expected profit on the option leg is, to first order, vega multiplied by the difference between the implied volatility received and the volatility subsequently realized. If those two are equal, expected profit on that leg is exactly zero. Call writing carries no independent edge; every dollar of excess return is the wedge. (§5)

2 · The premium is a correlation premium more than a volatility premium. Index options carry a large negative variance risk premium; options on the individual components carry little or none. The premium exists because investors overpay for basket protection against correlations spiking in a crisis. This is why selling calls on a high-volatility single name is not the richer trade it appears to be. (§3)

3 · Selection should screen on the wedge per unit of risk, not on implied volatility level. High implied volatility on a single stock signals risk that is priced approximately fairly. The mispricing sits in baskets. Screening on implied volatility level points systematically toward the worst risk-adjusted opportunities available. (§4)

4 · Harvest scales with the square root of roll frequency; costs scale linearly. Writing twelve monthly calls collects roughly the square root of twelve, about 3.5 times, what a single annual call collects, not twelve times. Because costs grow linearly while harvest grows as a square root, an interior optimum exists and has a closed form. (§5.3)

5 · The buffer raises the floor; it does not reduce the loss. Premium collected in one period does not stand ready to absorb a drawdown in a later one. What carries forward is realized cash compounding into a higher base, so the portfolio enters and exits a decline at a higher absolute level. A given percentage decline on a larger balance costs more dollars, not fewer. The benefit is the higher starting point, not a smaller loss. (§6)

6 · The richest premium arrives at the worst moment to capture it. Implied volatility spikes during market declines, so the fattest premiums coincide with the periods when capping the subsequent rebound is most costly. The correct response to a wide wedge is to move strikes further out, not to write more. (§7)

Three standing rules

1. **Write only on an underlying you want to hold through a full cycle.** The overlay is not a substitute for a view. If the position falls materially, accumulated premium will not save it.
2. **Write only where there is volatility to sell.** Insufficient implied volatility produces premium too thin to justify the capped upside or the operational burden.
3. **Write only where the options market is liquid enough to roll.** Bid-ask leakage on every roll compounds directly against a strategy whose entire excess return is a few volatility points.

What would change these conclusions

- A persistent collapse in the variance risk premium. The strategy's excess return is the wedge. If crowded selling compresses it toward zero, the overlay stops paying and there is no secondary source of return to fall back on.
- A client with genuine need for the right tail. For a portfolio funding a long-dated liability that requires equity-like compound growth, capping upside is a real cost that risk-adjusted metrics understate.
- Tax treatment. Written calls generate short-term gains and can affect the holding period of the underlying. In a taxable account this materially changes the arithmetic and is not modeled here.

◆

2 · WHAT YOU ARE ACTUALLY SELLING

A covered call is long stock plus a short call at strike K. Put-call parity, for European options on the same underlying, strike and expiry, states:

$$C - P = S \cdot e^{(-qT)} - K \cdot e^{(-rT)}$$

Taking the simplest case of a non-dividend-paying underlying (q = 0) and rearranging for the covered call position:

$$S - C = K \cdot e^{(-rT)} - P$$

The left side is the covered call. The right side is a risk-free bond paying the strike at expiry, plus a short put. Where the underlying pays dividends, the covered call additionally captures their present value, which adds a term of $S \cdot (1 - e^{(-qT)})$ but does not change the structure.

> A covered call is economically equivalent to holding a bond and selling a put. This is an identity, not an analogy.

The practical importance is presentational. A covered call program is routinely described as conservative and income-generating, while a put-writing program is described as speculative and risky. They are the same position. An investment committee comfortable with one and not the other is responding to labeling rather than to risk.

The identity also clarifies the payoff shape. A short put has capped gains, the most that can be earned is the premium, and a fully intact left tail. The distribution of outcomes is therefore negatively skewed by construction: many small positive results, occasional large negative ones. Every measurement decision in §8 follows from this.

None of this makes the strategy unattractive. Selling insurance is a legitimate and well-compensated business. But it should be recognized as what it is.

3 · WHERE THE PREMIUM COMES FROM

The strategy's excess return depends on implied volatility systematically exceeding subsequently realized volatility. That gap is the variance risk premium, and its structure is more specific than usually assumed.

Driessen, Maenhout and Vilkov found that index options carry a large negative variance risk premium while options on the individual index components do not exhibit one. Their evidence for the mechanism is direct: average implied correlation of 46.7% against realized correlation of 28.7%. Option markets price baskets as though constituents will move together substantially more than they actually do.

The economic explanation is straightforward. Investors buy index options as portfolio insurance, and what portfolio insurance must protect against is precisely the event where diversification fails and everything falls together. Demand for that protection is concentrated and relatively price-insensitive, so it is bid above fair value. Individual stock options attract no comparable systematic hedging demand.

Carr and Wu documented the resulting gradient: variance risk premia strongly negative for the S&P 500, S&P 100 and Dow Jones Industrial Average, but negative with smaller absolute magnitude for the Nasdaq 100 and for most individual stocks.

The two findings differ in degree, and the difference is worth acknowledging rather than smoothing over. Driessen and colleagues find individual variance risk essentially unpriced; Carr and Wu find it priced, but weakly. The samples and methods differ. What both support is the ordering, which is the part this memo relies on: the premium is largest in broad indices, smaller in narrower and more concentrated baskets, and smallest or absent in single names.

The reason is internal to the mechanism. The correlation premium is the gap between implied and realized correlation. In a narrow sector, constituents are already highly correlated in reality, leaving less gap to harvest. In a broad diversified index, realized correlation is low while implied correlation is priced high, and the gap is wide.

> More volatile underlyings pay more premium in absolute terms, but premium quality, the portion representing genuine mispricing rather than fair compensation for risk, is highest where the basket is broadest, not where volatility is highest.

4 · SELECTION: PREMIUM QUALITY VERSUS PREMIUM QUANTITY

The screening criterion that follows is implied volatility minus subsequently realized volatility, assessed per unit of risk borne, not implied volatility level.

An illustration of the structure, using round numbers for exposition rather than measured values:

Illustrative figures for exposition. Not measured values.

UNDERLYING	IMPLIED	REALIZED	WEDGE
Broad index	18%	15%	3 points
Sector fund	24%	22%	2 points
Single stock	45%	44%	1 point

On these illustrative figures the single name pays roughly two and a half times the implied volatility of the index and delivers roughly a third of the edge. The premium is larger because the risk is larger, and the compensation is approximately fair. Fair compensation is not an opportunity.

Two complications keep this from resolving into a simple preference for broad indices.

Absolute dollars still matter. A 2-point wedge on a 24-volatility underlying can produce more premium income than a 3-point wedge on an 18-volatility underlying, because vega scales with the underlying's price and the option's absolute premium scales with its volatility level. A client whose objective is distribution funding rather than risk-adjusted return may reasonably prefer the sector fund, and should do so knowingly.

Sector premia are heterogeneous. Published work using sector option data reports sector-specific implied correlations and risk premiums that vary widely across sectors. There is no stable ranking to memorize. Sector ETF options only became liquid enough for reliable estimation in the latter half of the 2000s, which further limits available history.

The honest conclusion is that this is a screening exercise to be run on live data, sector by sector, and refreshed over time, not a static preference. §10 specifies the screen.

◆

5 · THE ENGINE: VEGA TIMES THE WEDGE

5.1 The central result

Consider a call written at implied volatility σ_i where the underlying subsequently realizes σ_r. The expected profit on the option position is the difference between the price received and the option's value under the true data-generating process:

$$\Pi = C(\sigma_i) - C(\sigma_r)$$

For a modest wedge, a first-order Taylor expansion in volatility gives:

$$\Pi \approx \nu \cdot (\sigma_i - \sigma_r)$$

where ν is the option's vega, the derivative of option value with respect to volatility.

> Expected profit on the option leg equals vega multiplied by the volatility wedge.

A necessary boundary on this result. It isolates the option leg. A covered call also holds the underlying, whose expected return is a separate matter entirely and is not addressed by this expression. The claim is not that a covered call earns only the wedge; it earns the equity return plus the wedge, minus the value of the upside surrendered. The claim is narrower and more useful: the option-writing activity itself contributes nothing beyond the wedge. Everything else the position earns, it earns by being long equity, which requires no options program.

Three consequences follow.

If the wedge is zero, the option leg contributes zero. Setting $\sigma_i = \sigma_r$ gives $\Pi = 0$ exactly. This is not an artifact of the approximation; it is the no-arbitrage content of the Black-Scholes framework. Writing calls generates no expected excess return of its own.

This also imposes a modeling requirement, addressed in §10: any simulation must specify pricing volatility and path volatility as separate inputs. A simulation using a single volatility for both will show the option program earning nothing in excess, and will have demonstrated only that the model is internally consistent.

Vega peaks at the money. For an at-the-money option, vega is approximately $0.4 \cdot S \cdot \sqrt{T}$, declining as the strike moves away from spot. The strike decision is therefore an explicit tradeoff: at-the-money strikes capture the most wedge per contract and surrender the most upside; out-of-the-money strikes preserve upside and harvest proportionally less. This is a quantifiable frontier rather than a matter of taste.

5.2 Roll frequency scales as a square root

Because vega scales with $\sqrt{T}$, writing n options per year of tenor $1/n$ produces an annualized harvest proportional to $n \cdot \sqrt{(1/n)} = \sqrt{n}$.

Twelve monthly options collect approximately $\sqrt{12} \approx 3.46$ times what one annual option collects, not twelve times. Monthly writing beats quarterly by $\sqrt{3} \approx 1.73$, roughly 73% more harvest, not three times more.

This is worth stating explicitly, because the common intuition that shorter options collect theta more often is directionally right and quantitatively misleading by a wide margin.

5.3 The optimal roll frequency

Gross annualized harvest grows as $\sqrt{n}$. Trading costs, bid-ask spread and commission, incurred once per roll, grow as n. Net harvest is therefore:

$$Net(n) = 0.4 \cdot S \cdot \Delta\sigma \cdot \sqrt{n} - c \cdot n$$

where $\Delta\sigma$ is the wedge expressed in decimal volatility terms and c is the all-in cost per roll in the same currency units as S. Differentiating with respect to n and setting to zero gives:

$$n^* = (\, 0.2 \cdot S \cdot \Delta\sigma \,/\, c\,)^2$$

An interior optimum exists, and it has a closed form. Writing too frequently destroys the harvest through costs; writing too rarely forgoes the square-root benefit.

The optimum is highly sensitive to the cost assumption, which enters squared. This makes execution quality a first-order determinant of the strategy's viability rather than an implementation detail, and it is the analytical justification for standing rule three.

The expression should be read as a scaling relationship rather than a precise prescription. It assumes a constant wedge, constant cost per roll, and at-the-money strikes, none of which hold exactly. Its value is in showing that an optimum exists and what governs it.

6 · THE BUFFER RAISES THE FLOOR

A covered call program cushions declines, and premium accumulates over time. Both halves of that sentence are true, and the relationship between them is frequently misstated.

What does not carry forward. Premium collected in January does not stand ready to absorb a decline in September. Each period's premium offsets only that period's move. The cushion does not stack against a single future drawdown, and describing the strategy as building protection over time in that sense is incorrect.

What does carry forward. Premium is realized cash and it compounds into the balance. After several years of writing, the account sits at a higher level than it otherwise would, so the portfolio enters any subsequent decline from a higher point and exits it at a higher point.

This is a level effect, not a loss-reduction effect, and the distinction is not pedantic. A 20% decline on a balance of 110 costs 22, while the same decline on 100 costs 20. The larger balance loses more dollars. What it does is end at 88 rather than 80. The benefit is entirely in where you started, not in how much of the decline you avoided.

> The buffer raises the floor rather than reducing the loss, and it accrues to the balance rather than to the position.

The within-period asymmetry is structural. Premium accrues approximately linearly in time. Drawdowns arrive convexly: most of a severe decline occurs in a small fraction of the period. A program collecting 1 to 2% per quarter accumulates perhaps 4 to 8% of within-period cushion over a year, against a decline that may deliver 30% in a matter of weeks.

The 2008 experience illustrates the magnitude. The CBOE S&P 500 BuyWrite Index declined approximately 29% against roughly 37% for the S&P 500 on a total return basis, approximately eight points of relief in the worst equity year in modern memory.

That relief came from premium collected during 2008, not from premium banked in the years before it. Each month of the decline, the program was short a call whose premium offset part of that month's move. This is the clearest available demonstration that the cushion is a within-period mechanism repeated many times, not a reserve accumulated in advance. Eight points across a full year of writing, against a 37% decline, is a genuine and material improvement. It is not downside protection, and a memo that called it protection would deserve the objection it received.

The recommended exhibit is cumulative premium collected plotted against drawdown depth, which makes the case and pre-empts the objection in a single chart.

◆

7 · CONDITIONING ON THE WEDGE

The variance risk premium is not constant, and its variation appears partially forecastable, which converts a static program into a conditional one.

Volatility mean-reverts. Implied volatility on index options tends to exceed subsequently realized volatility, and elevated readings tend to be followed by declines toward long-run averages. When implied volatility is high, there is more room for realized volatility to come in below it, widening the captured wedge.

The wedge itself is the better signal. Bollerslev, Tauchen and Zhou found that the difference between implied and realized variance predicts stock market returns, explaining a non-trivial fraction of quarterly variation and dominating conventional predictors such as the price-earnings ratio and dividend yield. Related cross-asset work finds that implied variance running above recent historical variance forecasts a larger subsequent variance premium.

The signal is therefore implied minus recent realized, not the implied level alone. A VIX reading of 25 following a calm period is rich. A VIX reading of 25 while the market is genuinely moving 25 is fair, and writing into it earns nothing.

The rebound trap

Here the strategy meets a structural problem that a naive volatility-timing rule walks directly into.

Implied volatility spikes when markets fall. So the moment offering the fattest premium is also the moment when a subsequent rebound is most likely and most violent. A program that responds to a wide wedge by writing more aggressively will collect an elevated premium, perhaps 4%, and cap a recovery that may run 15% or more.

> The premium genuinely is richer. The opportunity cost rises at the same time, and faster.

The design rule that follows: when the wedge is wide, move strikes further out of the money rather than writing more contracts or writing closer to the money. This retains the elevated premium while restoring room for the recovery. Vega declines as the strike moves out, so some harvest is surrendered, but §5.1 quantifies exactly how much, making this a solvable optimization rather than a judgment call.

This single rule is, in practical terms, the difference between a covered call program that survives a market cycle and one that is abandoned after capping a recovery.

8 · MEASURING IT PROPERLY

Section 2 established that the payoff distribution is negatively skewed by construction. That has direct consequences for evaluation.

The Sharpe ratio systematically flatters short volatility strategies. Sharpe divides excess return by standard deviation, which treats upside and downside variation identically and assumes an approximately normal distribution. A covered call program deliberately removes upside variation, the truncated right tail, which lowers measured standard deviation without removing any downside risk. The ratio improves for a reason that has nothing to do with the strategy being better.

This is not a subtle effect. It is the mechanism by which a range of option-selling strategies have historically appeared attractive on Sharpe while carrying risks Sharpe cannot express.

Any honest evaluation should report, alongside Sharpe:

- **Sortino ratio** – penalizes downside deviation only, removing the credit given for eliminated upside
- **Skewness** – quantifies the asymmetry directly, and should be visibly more negative than the underlying
- **Maximum drawdown** – a path-dependent measure that captures what the buffer does and does not do
- **Conditional value at risk at the 5% level** – average loss in the worst outcomes, which is where a short volatility position's risk actually lives

The recommendation is not that Sharpe be discarded, but that the strategy be evaluated on a panel. Where Sharpe improves and Sortino does not, the improvement is an artifact of the metric.

9 · STANDING RULES

9.1 Write only on an underlying you want to hold through a full cycle

The overlay generates income and may improve risk-adjusted return. It does not rescue a position that falls materially. Accumulated premium is measured in single-digit percentages per year against declines that can run tens of percent. The strategy presumes the underlying is worth owning; it is not a substitute for that judgment, and it should never be the reason a position is retained.

9.2 Write only where there is volatility to sell

Insufficient implied volatility produces premium too thin to compensate for the capped upside. It also tends to coincide with thin option volume, so the strategy may be untradeable in practice even where quoted options exist.

9.3 Write only where the options market is liquid enough to roll

Section 5.3 makes this quantitative: cost per roll enters the optimal-frequency expression squared. A strategy earning a few volatility points cannot survive wide bid-ask spreads compounded over dozens of rolls. This rules out most single names and many thinly traded sector products regardless of how attractive their quoted premium appears.

10 · METHODOLOGY FOR THE DISTRIBUTIONAL ANALYSIS

The analytic results above are complete. The distributional results require simulation, specified here and pending execution.

Pricing. Black-Scholes, verified against put-call parity before use. Pricing volatility is set as trailing realized volatility plus an explicit wedge, toggled in volatility points. The wedge is the single most important input in the model.

Volatility estimation window. Matched to option tenor as the base case, a 30-day option priced off 30-day trailing realized volatility, consistent with the construction of standard variance risk premium measures. The window is exposed as a toggle. Short windows are statistically noisy, with a standard error on the order of $\sigma/\sqrt{(2n)}$, and volatility mean-reversion means longer or blended estimators such as the HAR specification forecast better. Matching is retained as the base case because real option markets exhibit the same recency bias, so the match reproduces market behavior rather than distorting it.

Skew. Implied volatility is reduced as strikes move out of the money, reflecting equity index skew. Pricing out-of-the-money calls at the at-the-money volatility would overstate premium received and bias results in favor of the strategy.

Path generation. Underlying paths are simulated at realized volatility, specified independently of pricing volatility. Student-t innovations are used in preference to Gaussian, so that drawdown depth and rebound violence, both material to §6 and §7, are not understated.

A note on look-ahead bias. The strategy's decision rules may use only information available at the time: trailing volatility, current spot, current implied. Realized volatility is never an input to a strike or timing decision. It is, however, a legitimate input to path generation, because a simulated path is a specified scenario rather than observed information. The distinction is between asking what happens if a wedge exists, which is valid, and allowing the strategy to know the wedge in advance, which is not.

Outputs. Full terminal wealth distributions for the overlay against buy-and-hold, across strike distances from at-the-money to 10% out of the money and tenors from monthly to quarterly. Reported on the panel specified in §8. The key deliverable is the breakeven wedge: the volatility points required before the overlay outperforms holding, by strike and tenor, net of costs. Outputs pending model execution.

11 · LIMITATIONS

Transaction costs are excluded from the base case and applied as an explicit toggled haircut, so their effect is visible rather than buried.

Taxes are not modeled. Written calls generate short-term gains and can affect the holding period of the underlying. In a taxable account this materially changes the arithmetic.

Black-Scholes assumes constant volatility. Real volatility is stochastic and clusters. The framework here captures the level effect that drives the result but not the dynamics.

A single wedge parameter stands in for a time-varying quantity. Section 7 establishes that the wedge moves and is partially forecastable. The base model treats it as constant, and the conditional strategy described in §7 is specified but not simulated.

The framework is a sensitivity study, not a backtest. It establishes how much premium quality is required for the strategy to justify itself under stated conditions. It does not claim those conditions obtained historically. A backtest is one path; this is a distribution, but it is a distribution conditional on assumptions, and should be read as such.

This analysis is a modeled framework prepared for illustrative purposes. It is not investment advice and does not constitute a recommendation of any strategy or security.