EUREX Architects of trusted markets

Research paper

Tail risk hedging in the credit market

Paying economy fares for first class safety

Tail risk hedging in the credit market – Paying economy fares for first class safety

Table of contents

Authors

Leon von Essen, Product Design, Eurex

Levin Schach, Client Solutions, Eurex

Summary

This paper serves as a playbook for navigating the turbulent world of high-yield credit, aiming to secure first-class safety while paying econ-omy-class fares. The strategy's foundation is that the market's view of the future, as observed through option prices, diverges from the historical distribution of actual market returns. We demonstrate that the options of the iShares Euro High Yield Corporate Bond UCITS ETF (IHYG) are consistently underpriced relative to the probability of severe drawdowns. This allows cheap tail-risk protection via deep out-of-the-money put options. This protection is then used as a buffer for a dynamically sized, leveraged long credit position. We show how the insurance premium can be covered by selling calls against the portfolio. A modest overwrite generates sufficient income to offset the cost of the protective puts, creating a nearly self-funded hedge.

The paper concludes by explaining the trade-off. While this approach provides a robust hedge, excessive overwriting can invert the portfolio's risk profile, turning a crash-protected strategy into a speculative short-volatility position.

Navigating turbulence – A first-class approach to high yield investing

In March 2020, iShares EUR high-yield ETF, IHYG, collapsed 24% within days as Covid news hit global markets. What initially felt like a gentle carry trade suddenly looked like a cliff edge. Credit portfolios that appeared diversified saw correlations spike, and many investors discovered they were on a turbulent flight without a seatbelt.

Downside protection via options can cushion such a crash. A put option acts like an insurance policy, which gains value as the underlying asset's price falls. During the March 2020 crash, as IHYG collapsed and volatility surged, the value of these protective puts exploded. The gains from these options directly offset the losses from the crashing ETF, effectively acting as the missing seatbelt in the turbulent market. This mechanism preserves capital when it is most at risk.

The key is to secure this protection without incurring excessive costs. In the credit markets, this can be achieved by comparing how future risk is priced in the options market with how risk has historically materialized. Option prices define a risk-neutral distribution (RND) of future returns via today's volatility surface, whereas realized returns follow a historical (physical) distribution (HD). The former is a pricing measure; the latter reflects realized probabilities.

Empirically, for IHYG options over 2019–2025, the left tail of the RND is systematically lighter than that of the HD. When the option-implied distribution underweights severe drawdowns relative to the historical distribution, deep out-of-the-money IHYG puts offer cost-effective crash protection. We aim to harvest this by pairing a monthly-rolled put hedge with a dynamically sized long credit exposure, designed to keep portfolio drawdowns within a pre-defined limit.

Backtests from 2019–2025 deliver superior compounding with capped losses, as drawdowns stayed below the target, while the Sharpe ratio increased significantly versus an unhedged IHYG allocation.

The following paper unpacks the different densities and the mismatch, outlines the trading mechanics, and compares several portfolio constructions. The result is a playbook for paying economy fare while enjoying first-class safety when turbulence hits the credit markets.

Historical density – The market's diary

Imagine driving across town. Your GPS promises a 15-minute trip. Helpful, yet years of rush-hour experiences and surprise roadwork have taught you that the same journey can stretch to 40 minutes. Markets face a similar issue. The RND implied by option prices resembles the route suggested by your navigation app. The HD is your mental logbook that records every traffic jam you've actually hit on that route. Only the latter captures how much journey times can vary.

Why deal with the past when all we care about is the future?

We invest our money in the physical world, so we earn and lose money under physical probabilities, not the risk-neutral ones used for pricing. Expected P&L is driven by how often losses occur in the real world. If tail events occur more often than option prices imply, crash insurance is effectively on sale. The HD is not just a naive forecast; it tells us how widely outcomes can spread and thus shows how heavy the tail was when credit markets turned. Used this way, the HD complements the RND. When both maps disagree, we can structure a hedge that is cheap under the RND but valuable under the HD.

But what is a density anyway?

Think of a density as a probability heatmap of future returns. The horizontal axis spans every τ-day return IHYG might realize. The vertical axis shows the likelihood of each outcome, based on past data. Peaks mark everyday moves, while tails warn of rare but violent crashes. The historical density in Figure 1 is built from past return windows and smoothed into a continuous curve. It shows how the market has evolved over time intervals of the same length.

Calculating the historical density without drowning in formulas (yet)

We fix a horizon τ (here, 46 trading days) and compute overlapping τ-day log returns for IHYG across the full sample. Conceptually, you can think of applying each historical τ-day return to today's price to create a cloud of possible future prices. These are then smoothed with a kernel density estimator to form a continuous curve, the historical probability density. The formal derivation and bandwidth choice are set out in the appendix.

Figure 1: The historical density function of log returns for the iShares Euro High Yield Corporate Bond UCITS ETF (IHYG) as of 2 December 2024. Returns are calculated over a 46-day period (τ = 46 days).

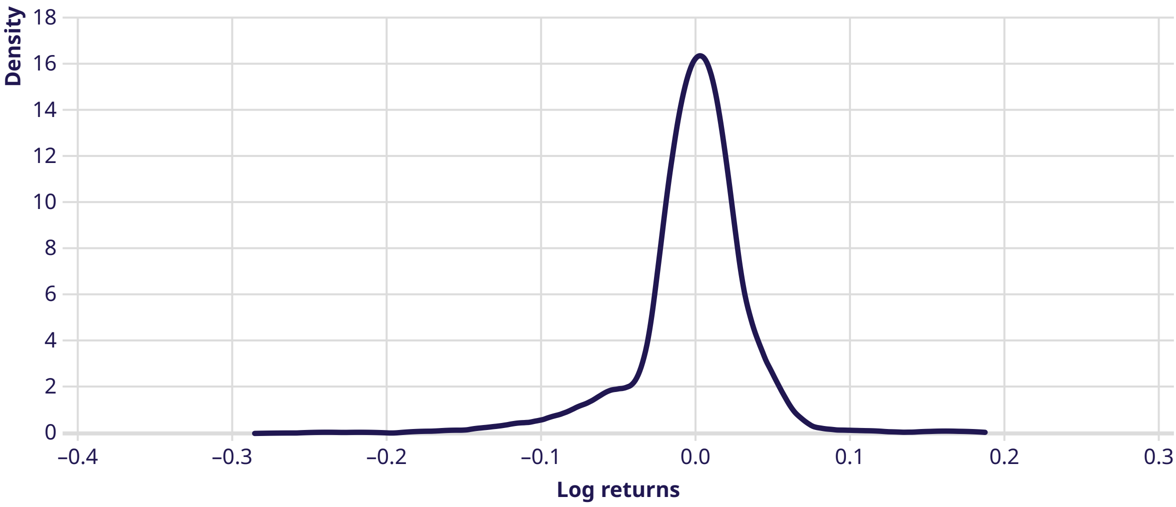

Risk-neutral density – the navigation app's probability map

Return to the navigation app metaphor. It promises a 15-minute drive, yet its code maintains a full probability distribution that perhaps gives a 70% chance of arriving in 15 minutes, a 25% chance of arriving in 18, and a 5% tail beyond 25 minutes. That curve is the software's model of traffic conditions.

Options markets have a similar curve called risk-neutral density. It reflects the risk-neutral probabilities of future returns implied by today's option prices. So, if you discount expected payoffs under this density at the risk-free rate, you recover the observed option prices.

But what does "risk-neutral" mean, and why is this pricing considered "fair"?

Let's take a step back and consider some basic derivatives theory. As with all derivatives, the option value must match the expected discounted cost of implementing a hedging strategy. Options are a little more complicated than linear instruments like futures because they require dynamic trading of the underlying and money-market instruments. "Risk-neutral" pricing refers exactly to this. The option price equals the discounted expected payoff under the risk-neutral measure, so that an option combined with its offsetting hedging strategy will earn the risk-free money-market return.

We can triangulate the map by using option prices across various strikes. Each call price at a given strike is the discounted expected payoff of that contract under the risk-neutral distribution. So, individually, each price contains some information about the return distribution. Taken together, the cross-section of option prices across many strikes paints a much clearer picture; they effectively pin down the entire risk-neutral distribution. The difference in prices between two neighboring strikes corresponds to the probability-weighted payoff between those strikes. The closer the strikes, the more precise the density for the area between them can be estimated.

Calculating the risk-neutral density without drowning in formulas (yet)

For a given expiry, we take the discrete set of call prices across strikes and fit a smooth curve using local polynomial regression. The second derivative of that fitted call price curve with respect to the strike is the RND itself. Full derivation and implementation details are provided in the appendix. The important point is that the RND is the option market's probability map.

The next step is to overlay the risk-neutral density on top of the historical density, as we discussed in the previous chapter. Where the RND's left tail sits below the HD's, crash insurance is selling cheap.

Figure 2: The risk-neutral density function of log returns for the iShares Euro High Yield Corporate Bond UCITS ETF (IHYG) as of 2 December 2024. Returns are calculated over a 46-day period (τ = 46 days).

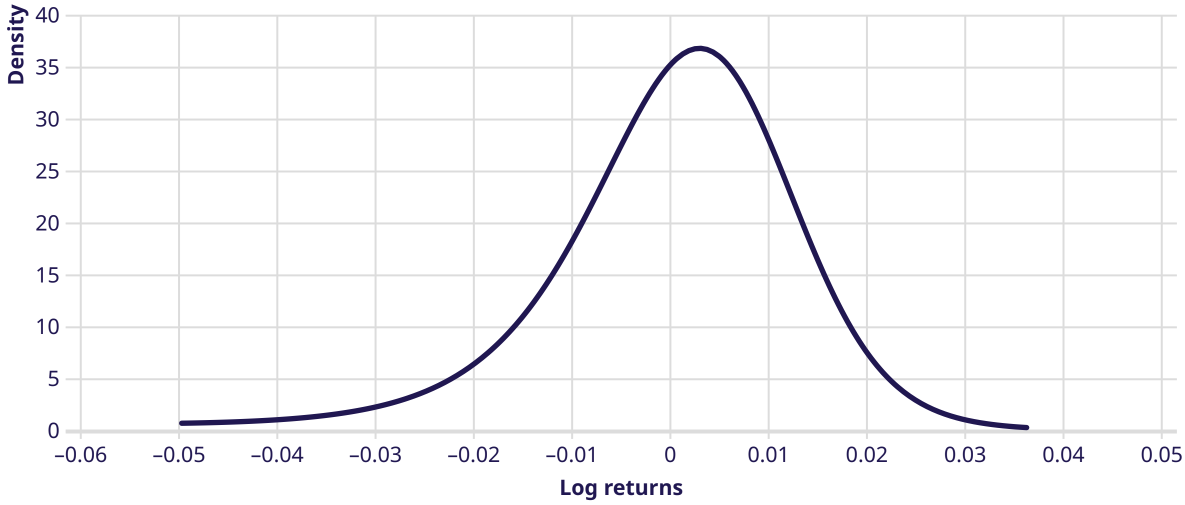

Putting the densities together – Can we pay economy fare for first class safety?

Figure 3 overlays the historical density (solid blue) and the risk-neutral density (solid green) for the same 46-day horizon. Two features are central to our approach:

1. Taller peak, thinner tails of the RND.
The RND (green) is both narrower and more sharply centered. Option prices therefore believe that most 46-day moves will be small and slightly positive.

2. Fatter left tail of the HD.
From around –3 % onward, the HD (blue) lies above the RND. In our sample, the historical data assigns roughly 2–3 × the probability of a –5 % or worse 46-day drawdown than the option market currently does.

This divergence is precisely what our hedge aims to capture. Deep-OTM IHYG puts are priced off the thin RND yet protect against the fatter HD. Relative to the historical distribution, the option market is effectively selling us crash insurance at a discount, providing a cost-effective buffer against portfolio drawdowns.

Figure 3: Historical (blue) and risk-neutral (green) probability distributions for 46-day log returns as of 2 December 2024.

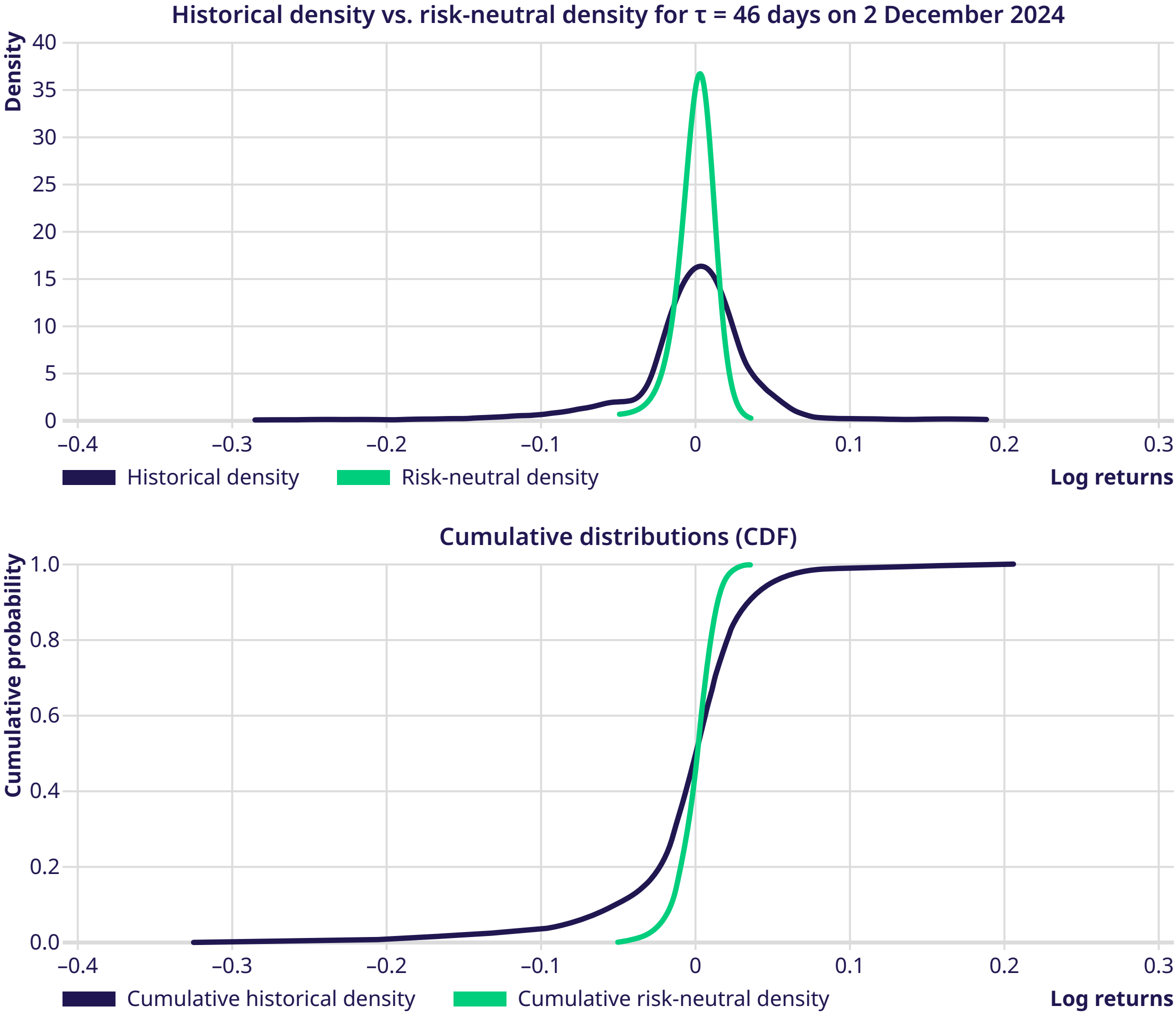

Theoretical framework – How the hedge adjusts

The strategy is built around deep out-of-the-money, 5-delta IHYG put options with an initial maturity of roughly 60 days, rolled monthly. At each roll date, we resize the portfolio's leveraged credit beta leg and the hedge notional to roughly maintain the total portfolio drawdown limit.

$$S_{Spot} - S_{Strike} \times Leverage = D_{Max}$$

The key mechanism is this dynamic sizing. The lower the put's strike (the further out-of-the-money the protection), the less leverage we can run while still respecting the same drawdown constraint. Once IHYG breaches the strike, delta approaches one, and the put's intrinsic value rises point-for-point with the loss on the long leg.

Each month, leverage, strike, and the number of put contracts are resized to reflect current conditions. When volatility is low and drawdowns are unlikely, leverage increases; when risk picks up, exposure is tightened. This structure systematically clips left tail risk without overdosing on theta.

Trading the leveraged high-yield protected portfolio

Two levers control the backtest's performance: the put's delta, which, in a sense, defines how far the strike sits below spot, and a maximum drawdown budget, which jointly limit the leverage applied to the underlying total return index. A 36 × 16 parameter grid, resulting in 576 portfolios, highlights how sensitive the performance is to these choices. The heatmap of cumulative ln returns shows that the highest-performing portfolios cluster at the riskier end of the drawdown parameter.

The global optimum in terms of maximum returns over the 2019 to 2025 window is a 40% drawdown limit paired with a delta of –0.05. This compounds to a cumulative return of about 1.08 and comes at the cost of aggressive exposure, as leverage stays at the cap of 6x for most months, except during the two drawdown periods. Annualized volatility is 0.36, putting the Sharpe at roughly 0.50.

With a 5% drawdown limit and a −0.05 delta put, the risk engine keeps leverage near 1× for most of the sample. The portfolio is essentially unlevered except during stress events. During major sell-offs, the tail hedge pays out, but for most of the time, the portfolio simply tracks the underlying IHYG total-return index, resulting in the lowest total return in the grid.

From the perspective of optimizing risk-adjusted returns, this set is almost optimal. Slightly tweaking the puts' delta to –0.07 maximizes the Sharpe value at 0.73. A very cheap tail-risk hedge is rolled every month. It does not erode portfolio gains in normal markets but protects against black swan events. The underlying position is only modestly levered. For most of the sample, no borrowing is required, and after mid-2023, average leverage is about 1.3×, with a peak of 2.4×. Overall, this combination delivers a cumulative log return of 0.4.

A compelling middle ground is a 15% max drawdown with a −0.05 target put delta. This combination delivers a Sharpe ratio of about 0.66 and a cumulative log return of 0.72, roughly doubling capital over the sample. The portfolio's effective leverage (Figure 5) averages around 1–3× after the Covid crash and 2–6× during the post-2023 relief period. In calm markets, the put strike hovers around 95% of spot; in turmoil, it pulls back toward 82.5%.

Figure 4: Heatmaps of cumulative log return (left) and Sharpe ratio (right) as a function of the Max Drawdown (x-axis) and Target Delta (y-axis) parameters.

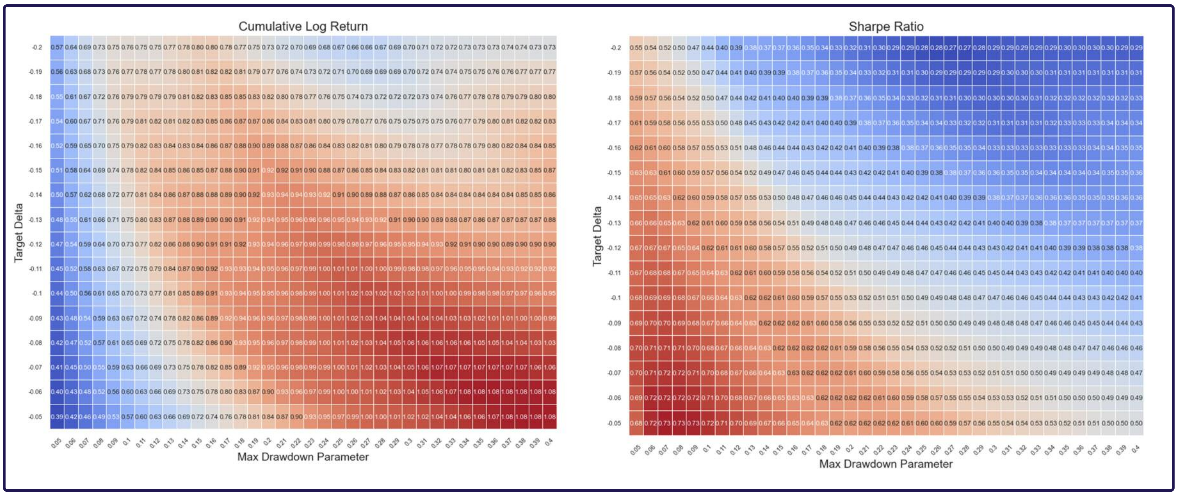

Figure 5: Time series of the traded put options' moneyness (blue) and the portfolio's beta (green) since 2020.

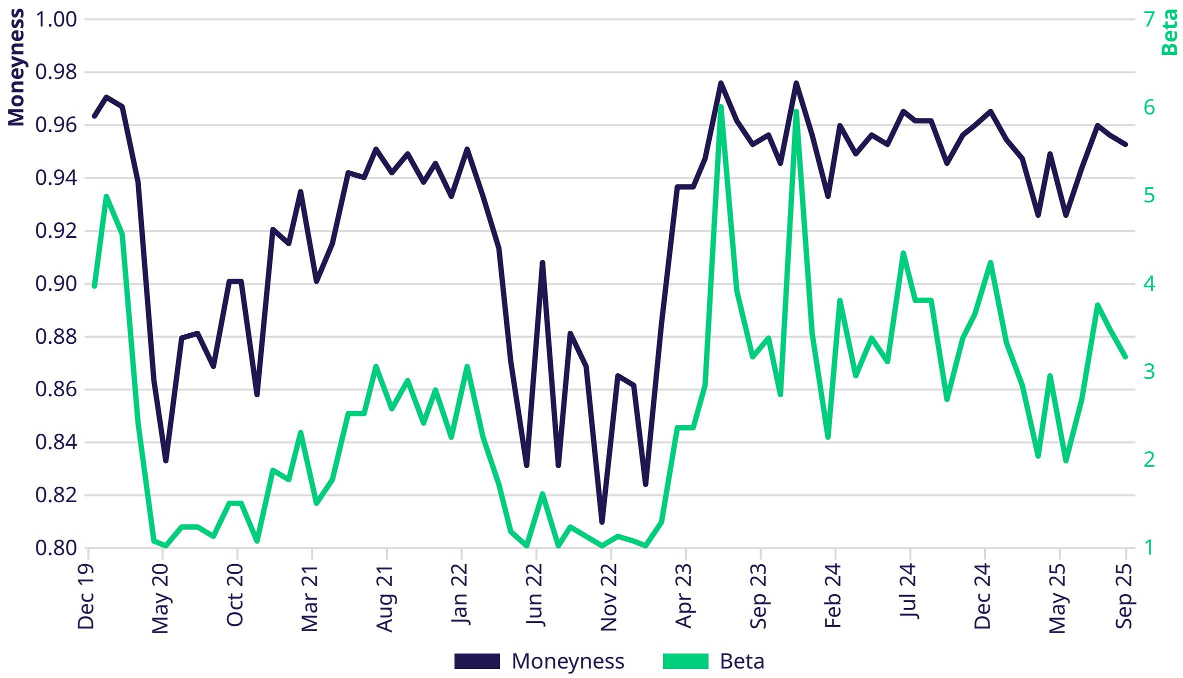

Within this 15% drawdown budget, we also test deeper hedges with target deltas from −0.10 down to −0.20. While all of these portfolios respect the same risk budget, effective leverage varies widely because closer-to-the-money puts allow more long exposure for a given drawdown constraint. The most conservative member of this set (delta = –0.05), which we just discussed, has a Sharpe value near 0.66, while the most aggressive variation (delta = –0.20) returned roughly 0.80 in gain, with 0.35 volatility and a Sharpe around 0.37, all within the 15% loss budget.

Figure 6: Performance of four portfolios with the same maximum drawdown and different target deltas, compared to the total return index. The portfolio deltas are –0.05 (dark blue), –0.10 (grey), –0.15 (green), and –0.20 (light blue).

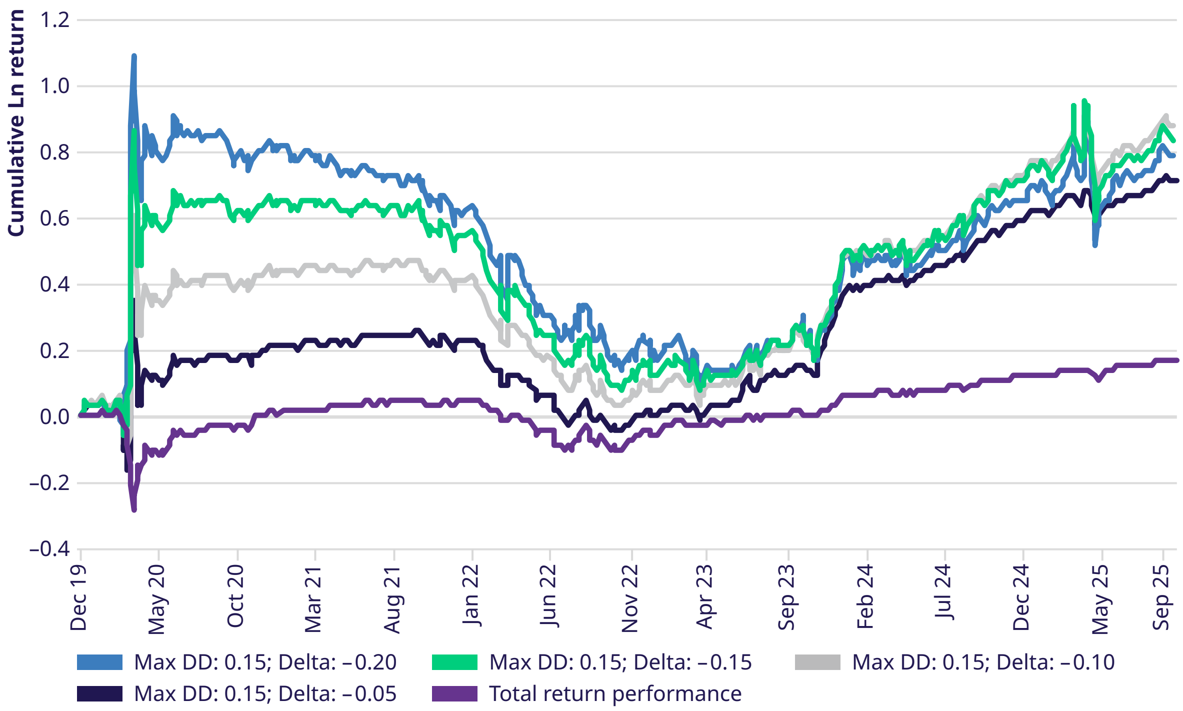

Figure 7: Leverage taken in the four sample portfolios with the same maximum drawdown and different target deltas. The portfolio deltas are –0.05 (dark blue), –0.10 (grey), –0.15 (green), and –0.20 (light blue).

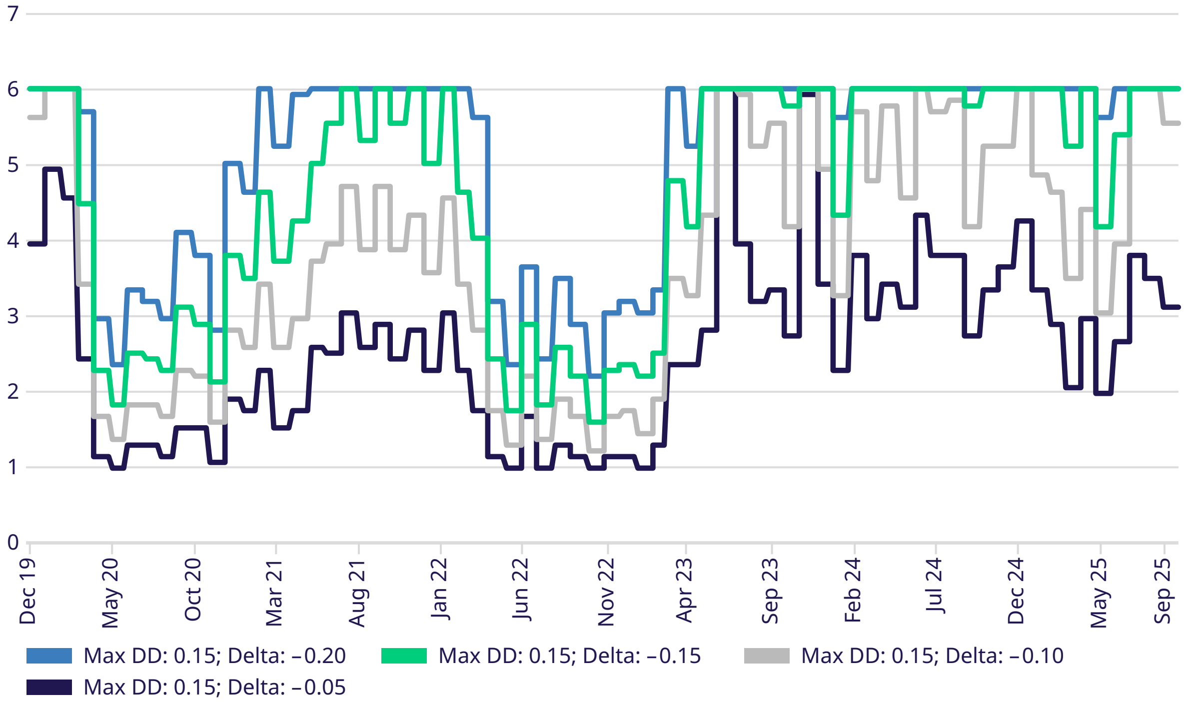

Taking a closer look at the 15% DD and –0.05 delta combination, the five-year path breaks naturally into four market regimes:

Table 1: Regime-specific performance; max DD 15%, delta –0.05

	ΔNAV	Return (%)	Running max DD	Ann. volatility	Options P&L	VaR	CVaR	Leverage range
Covid- Crash	36,851€	4%	–27%	64%	387,910€	–5%	–8%	2 – 5
Recovery 1	207,415€	20%	–6%	8%	–72,949€	–0.7%	–1%	1 – 3
Russia-Ukraine-War	–263,509€	–21%	–24%	12%	–1,727€	–1%	–2%	1 – 3
Recovery 2	1,036,766€	103%	–13%	9%	–117,210€	–0.6%	–1%	1 – 6

Covid Crash (March 2020)

IHYG's total return index collapsed from 254 to 207 (–23 %) in just 15 trading days. The portfolio entered the sell-off at roughly 2.4× leverage, protected by puts 6% below spot. Pre-crash, the puts declined as the market had a final upside breath, but as headlines around the world on 11 March broke, implied vol exploded and put value surged ~21×. When IHYG bottomed at 207, the hedge had returned 198×, offsetting losses from the levered long. The marked bounce before the scheduled roll date erased a big chunk of that unrealized gain, but the strategy still outperformed buy-and-hold materially.

First Recovery (mid-2020 – end-2021)

Volatility normalized, and the position was progressively re-levered from ~1× to ~3×. Most put series expired worthless, and theta became the dominant cost of insurance, while the underlying leg drove the cumulative log return to about 0.24.

Rate hikes driven by Slow Bleed (Q1 2022 – Q4 2022)

Russia's invasion of Ukraine introduced a persistent grind lower in credit without the kind of crash dynamics that activate tail hedges. Option premiums again decayed to zero, but this time the leveraged long leg provided no upside to compensate, producing a severe drawdown period.

Second Recovery (2023 – 2025)

Once rate hikes stopped and economic uncertainty receded, realized volatility fell. The risk engine turned up the leverage, lifting the portfolio from a cumulative log return of −0.04 back to +0.72 in roughly two and a half years.

This underscores a classic convexity trade-off. Large drawdown allowances purchase upside beta at the cost of severe theta bleeding, requiring a high drawdown tolerance. Tight limits preserve capital but sacrifice compounding.

Only a handful of put series generated a positive P&L relative to their gross notional exposure. This is the essence of crash insurance. You bleed modest premiums for many years to be protected during the very rare, very large dislocations.

Figure 8: Put P&L in % of gross notional exposure of each put option series over its lifespan.

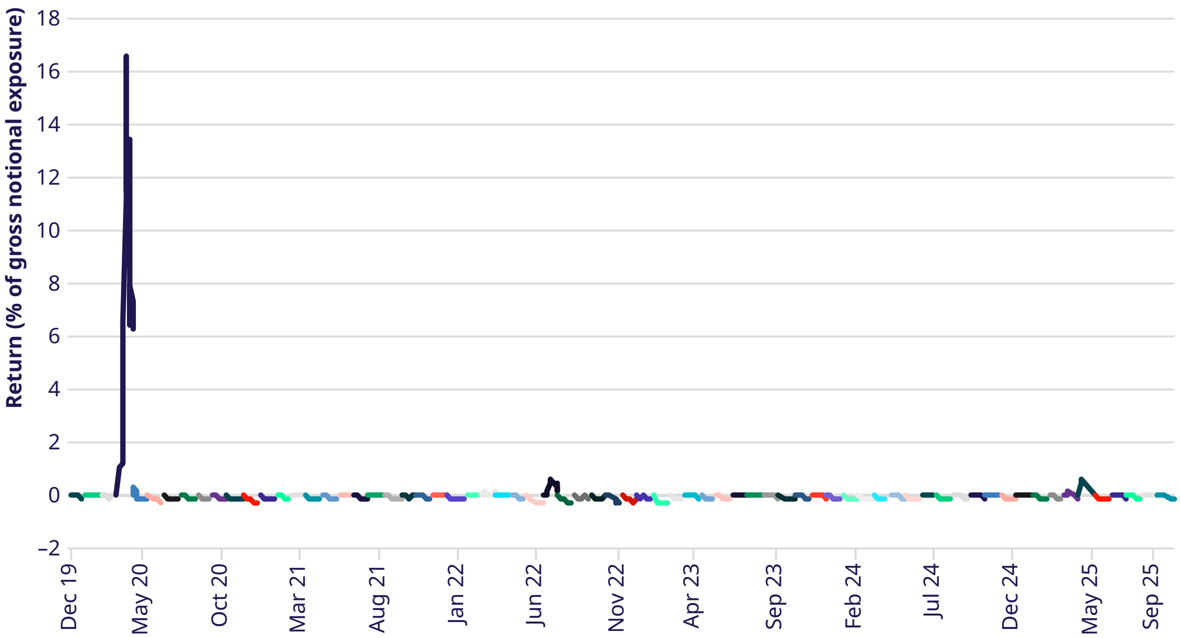

A call a day keeps the theta away

In the previous chapter, we showed that a dynamically levered long position in IHYG, hedged with a strip of deep out-of-the-money 5-delta puts, can compound attractively while respecting a hard 15% drawdown limit. One drawback of this strategy is the constant theta bleed. The portfolio pays an insurance premium every month to keep large drawdowns off the table.

The next step is to ask whether this insurance can be, at least partially, self-financed. Rather than funding the put leg entirely from the underlying risk premium, we can lease out a small slice of the portfolio's upside. Concretely, at each monthly roll date, we sell 35-delta calls on a fraction of the IHYG total-return notional the portfolio already owns. This creates a covered overwrite, as we are long at least one unit of IHYG total return for each call. We never sell naked calls, and for modest overwrites, the portfolio remains net long IHYG. What changes is the distribution of outcomes: we trade away a piece of future right-tail convexity in exchange for immediate option premium that pays the put leg bill.

At a given roll date, let N_{TR} be the number of IHYG total return units held and S the index price. We denote by m the call multiplier, meaning we sell $N_C = m \times N_{TR}$ calls at 35-delta, for a notional of $mN_{TR}S$.

In practice, we find that relatively light 35-delta overwrites with m ≈ 0.45 generate a monthly call income that, on average, covers 100% of the put leg's initial cost. So, a modest overwrite is already sufficient to move the tail hedge close to being self-financed. Most of the crash protection is funded by selling out-of-the-money upside.

Over the full 2019-to-2025 backtest, the overwrite with a call multiplier of 0.45 delivered a cumulative ln return of 0.66 at roughly 16% volatility, slightly underperforming the uncovered strategy. This illustrates that 'self-financed' does not equal 'free'. The hedge is funded by sacrificing upside, so total return can be slightly lower even though the explicit option outlay is neutral on average.

Figure 9: Performance of four portfolios with the same maximum drawdown of 15%, –0.05 delta puts and different call multiplier. The call multipliers are 0.2 (light blue), 0.5 (dark blue), 1 (green), and 0 (grey).

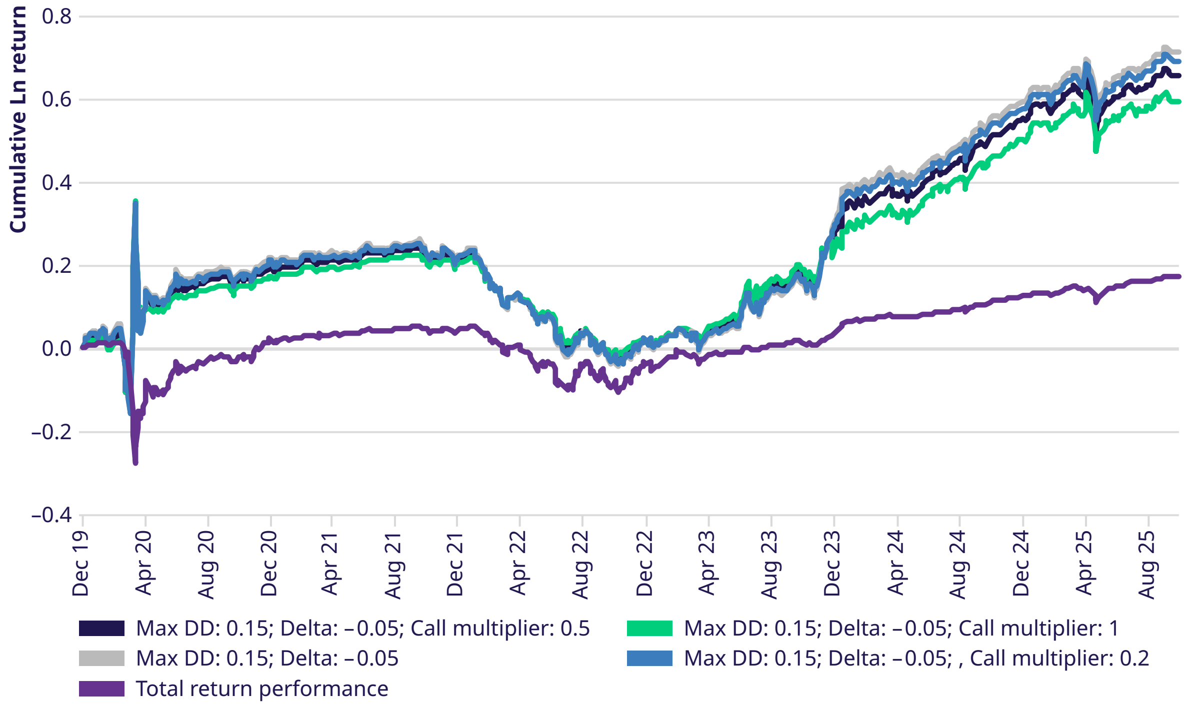

Putting on the Greek(s) goggles shows that for a small *m*, the call leg trims the portfolio's net delta moderately. Let Δ_{TR} be the delta of the total return exposure (+1), Δ_P the delta of the deep OTM puts (-0.05) and Δ_C = 0.35 the delta of the calls. For a given call multiplier m, the net delta of the overlayed strategy is equal to:

$$\Delta_{net} \approx N_{TR}\Delta_{TR} + N_P\,\Delta_P + (-N_C\,\Delta_C)$$

For small *m*, the strategy remains strongly long credit beta, with slightly lower upside convexity than the pure put hedged version. At the same time, the overwrite introduces a short-gamma, short-vega, long-theta component. We effectively sell convexity for carry.

The call leg adds negative convexity for large up-moves (and gamma). Combined with the long gamma of the deep puts, the net convexity profile becomes more S-shaped. This means strongly long convexity in the left tail, negative convexity in the right tail, but a net short gamma position at trade entry.

Deep OTM puts are long vega, while short 35-delta calls are short vega. A light overwrite brings net vega closer to zero, making the portfolio less sensitive to volatility spikes. However, all the options used are traded with less than two months to expiry and have limited vega to begin with.

The short convexity and vega position introduce carry generated via the time decay of the options. The put leg is strongly short theta, while the call leg is long theta. For moderate m, the net theta increases, meaning that most of the insurance bleed is effectively offset.

Figure 10 shows the call-premium-yield (CPY), the upfront cash from selling a 35-delta call, expressed as a percentage of the IHYG total-return notional, and the index's 21-day realized volatility. CPY is the monthly "rent" we earn for leasing out a slice of upside. From 2019-2025, the two lines rise and fall together. The Pearson correlation is $\rho \approx +0.6$, so the noisiest months are also the months in which we collect the most rent.

Why? When volatility increases, call prices increase due to their vega. Our short-gamma overwrite acts as a cash generator. Just as the 5-delta put hedge becomes costlier to carry, call income surges, helping pay the bill. In calm markets, the generator slows down but never shuts down. Premium still trickles in, slowing the relentless theta bleed.

Figure 10: Monthly call-premium-yield (CPY) and 21-day realized volatility for the IHYG index (2020 – 2025)

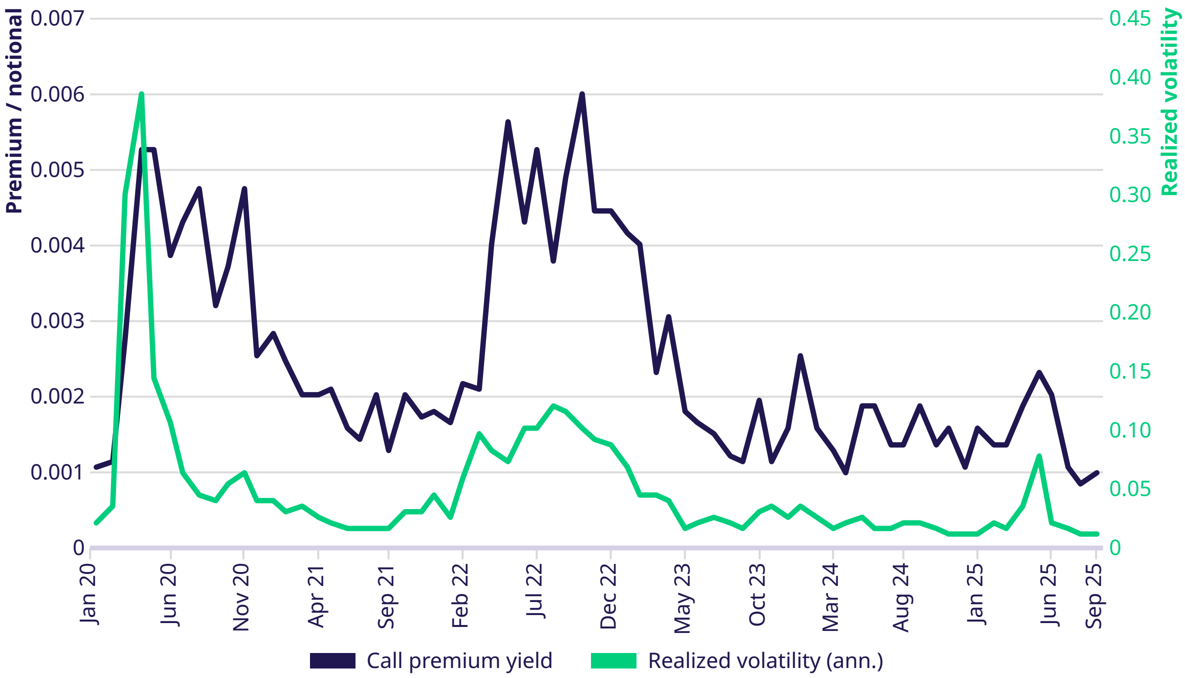

The catch appears after roll day. Put theta scales roughly linearly with volatility, so if volatility jumps mid-month, the daily bleed jumps with it. The overwrite premium, however, is locked in on day 0. Therefore, a volatility shock on day 10 leaves us with a funding gap even though the put's mark-to-market is surging. Short gamma is the price we pay, a steady income stream traded for some upside drag in violent rallies.

To summarize, the overwrite finances the tail hedge and slightly reshapes the right tail, while preserving the core characteristics of the strategy, namely levered exposure to credit beta with crash protection intact.

How far can we push self-funding?

The picture changes once the overwrite intensity is pushed beyond the light overlay theme. Consider a 3× overwrite. At each roll date, instead of selling calls on a small fraction of the IHYG total-return notional, we now sell calls three times the notional of the long IHYG exposure. From a risk perspective, the book is now heavily overwritten, as the short call position scales to a notional that exceeds the long side. In Greek terms, the net delta right after the roll, neglecting the puts, is approximately

$$\Delta_{net} \approx 1 \times N_{TR} - 3 \times N_{TR} \times 0.35 = -0.05 \times N_{TR}$$

The deep OTM puts add a small negative delta on top, reinforcing the short bias. So, while the strategy looks long IHYG on the surface, the net delta just after the roll is already close to zero or even negative.

Now consider what happens when the underlying experiences a sharp upward move over the month. As the index rallies toward the call strike, the delta of each call increases toward 1. The call leg's delta contribution transitions from roughly $-1.05 N_{TR}$ to roughly $-3N_{TR}$ when the calls are deep ITM. In the right tail, the overall portfolio has a $\Delta_{net} \approx -2N_{TR}$. After a large rally, the portfolio effectively becomes short 2 units of IHYG for every 1 unit held long. The P&L profile in the right tail is then dominated by the short call, not by the long underlying.

This explains the counterintuitive result in the backtest. Over a month in which IHYG jumps sharply higher, the portfolio can lose substantial value even though it holds a long IHYG position.

Figure 11: Cumulative log returns of a -0.05 delta put, levered long credit portfolio with a call multiplier of 3.

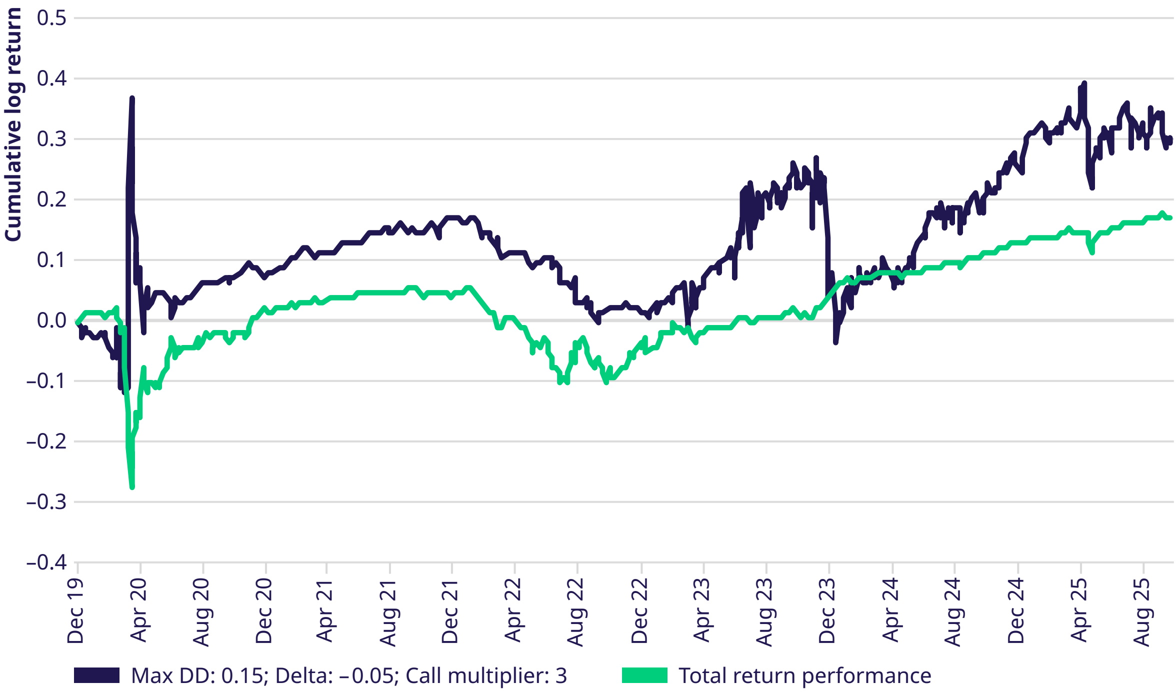

In addition to the pure option P&L, the portfolio must also post additional variation margin as the short calls go deep in-the-money. In the simulation, this shows up as substantial cash outflows into margin, potentially pushing the cash account negative and increasing reliance on funding at the borrowing rate. This does not change the total mark-to-market P&L; it is just a transfer between cash and margin.

In this regime, the overwrite is no longer a harmless way to finance the tail hedge by shaving off a small portion of upside. It fundamentally inverts the risk profile. On large monthly rallies, the portfolio behaves more like a leveraged short call position with a residual long underlying, rather than a long credit exposure with a crash hedge.

About Eurex

Eurex stands for the leading European derivatives exchange and – with Eurex Clearing – one of the leading central counterparties globally. Being architects of trusted markets characterized by market liquidity, efficiency and integrity, we provide our customers with innovative solutions to seamlessly manage risk. On the trading side, we mastermind the most efficient derivatives landscape by pioneering ingenious products and infrastructures as well as by building 'smart' into technology – offering a global product range, operating the most liquid fixed income markets in Europe and featuring open and low-cost electronic access. As central counterparty, Eurex Clearing builds trusted relationships with and amongst market participants, enabling effective risk management and delivering high efficiencies to clients.

Architects of trusted markets

Methodology

Filling the gap when an option is not quoted – SABR[1]

In every monthly roll, the strategy buys a two-month 5-delta put and sells a 35-delta call while going leveraged long the IHYG total return future. Unfortunately, far out-of-the-money strikes were only introduced in 2025 for IHYG, and for much of the observed period, the options used did not exist. When the strike is missing, we calculate the price in three steps following the SABR model.

1. Construct the forward price via put-call-parity

For each expiry T, we pick a call $C(K,T)$ and put $P(K,T)$ at the same strike K to calculate the synthetic forward. The put-call-parity implies:

$$C(K,T) - P(K,T) = e^{-r_T T}(F_T - K)$$

Where F_T is the forward value of the ETF and r_T the risk-free rate for maturity T. Rearranging gives:

$$F_T = \frac{C(K,T) - P(K,T)}{e^{-r_T T}} + K$$

As daily money market quotes only exist for a handful of tenors, the term structure r_T is interpolated with a cubic spline.

2. Building a continuous SABR smile

The scattered implied vols from the available quotes are not smooth and cover only a part of the delta range. To price an unquoted 5-delta put, we fit a single, continuous SABR smile by minimizing squared errors between market vols and the standard SABR implied-volatility approximation across strikes. The calibration is per maturity T.

$$\min_{\theta_T = (\alpha,\beta,\rho,\nu)} \sum_i [\sigma_{SABR}(F_T, K_i, \theta_T) - \sigma_{market}(K_i, T)]^2$$

[1] Hagan, Patrick & Kumar, Deep & Lesniewski, Andrew & Woodward, Diana. (2002). Managing Smile Risk. Wilmott Magazine. 1. 84–108.

3. Locating the X-delta Strike

Given the calibrated θ_T^* and F_T, we solve for the strike K, such that the forward put-delta equals the target (e.g., –0.05), using Black-Scholes and a root finding algorithm:

$$\Delta_{put} = N(d_1) - 1, \qquad d_1 = \frac{\ln\left(\frac{F}{K}\right) + \frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}}$$

Plug $K_{0.05}$ and $\sigma_{SABR}(F, K_{0.05})$ back into the Black-Scholes forward formula to get today's price.

4. Daily mark-to-market with updated SABR calibration

After the trade is set, we re-estimate SABR parameters each trading day from fresh quotes and revalue all option legs under the updated smile.

Figure 12: SABR volatility curve and prices on 2022-06-09 for the expiry 2022-12-16

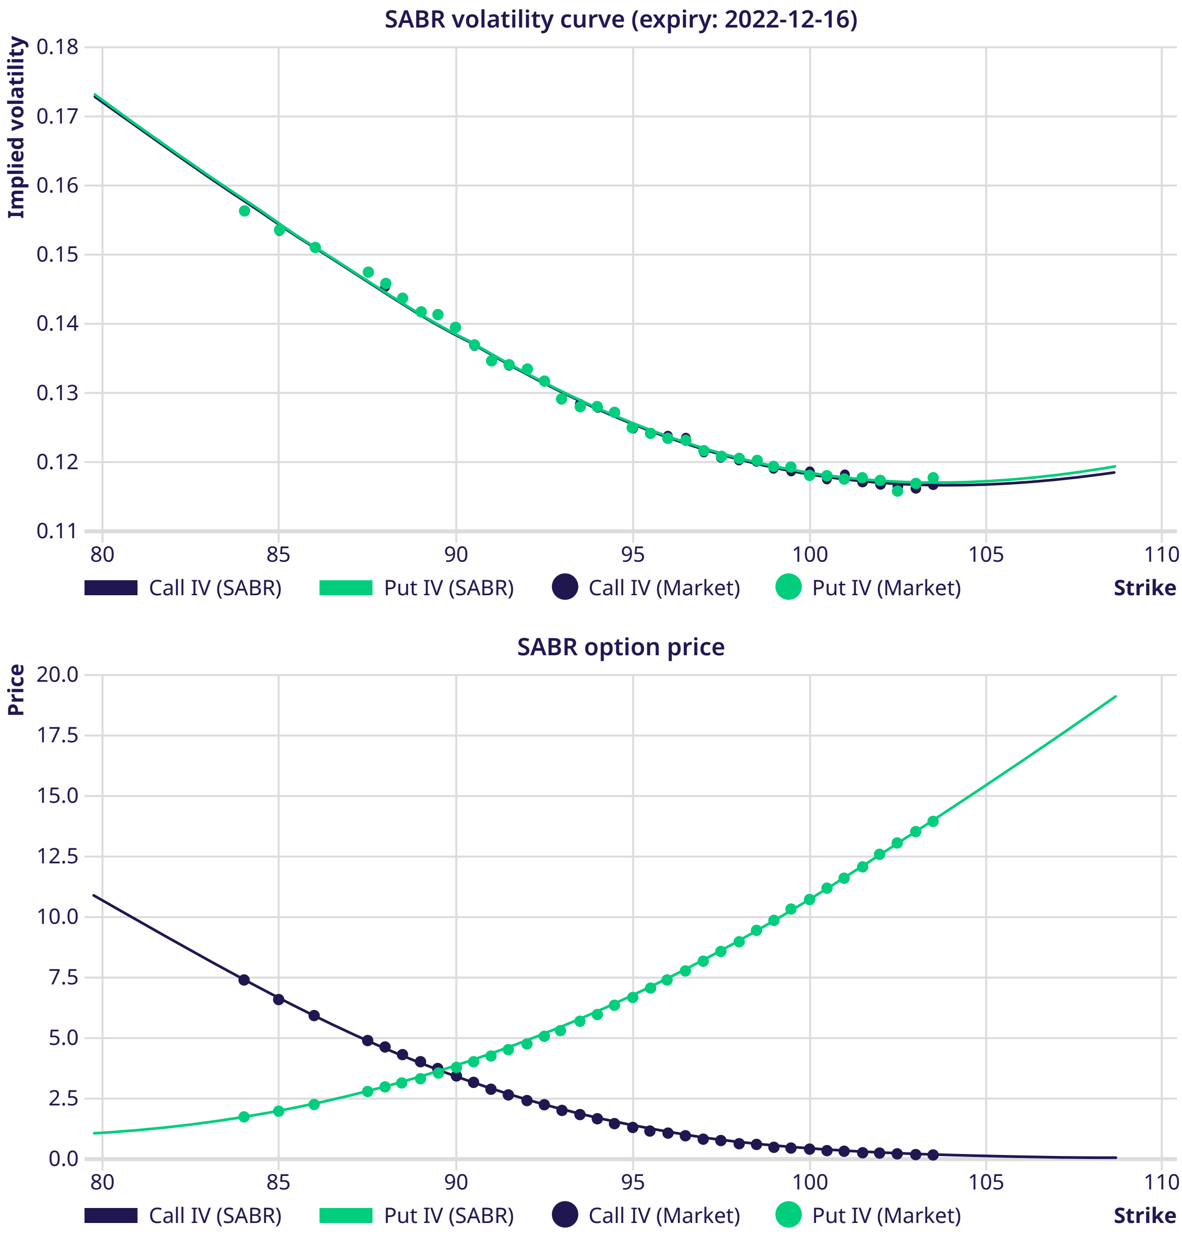

Historical density

To understand how risky IHYG has been historically under the physical measure, we consider the historical density, also known as the physical distribution. It reflects the empirical probability of future price moves based on actual past returns. It assumes that past price patterns carry information about future price movements and their probabilities.

To estimate the HD, we follow the nonparametric kernel-based method described by Duan et al. (2012)[2]. This uses a rolling window of logarithmic returns. Specifically, for each window of size τ, we compute:

$$r_\tau^k = log\left(\frac{S_{t-k}}{S_{t-k-\tau}}\right), \qquad k = 0, 1, 2, \dots J$$

where S_t is the spot price at time t, and r_τ^k represents the k-th log return in the backward-looking window. Each of these returns is then forward-projected to simulate potential future price levels at T:

$$S_T^k = S_t \times exp(r_\tau^k)$$

This produces a discrete set of future price scenarios $\{S_T^k\}_{k=1}^J$, which we transform into a continuous density estimate using kernel density estimation. The estimated historical density $\hat{p}_h(S_T)$ is given by:

$$\hat{p}_h(S_T) = \frac{1}{Jh}\sum_{k=1}^{J} K\left(\frac{S_T - S_T^k}{h}\right)$$

K(·) is a kernel function. We use the Epanechnikov kernel, and h is the bandwidth parameter that controls the smoothness of the estimate. A detailed discussion of kernel and bandwidth selection can also be found in Duan et al. (2012).

[2] DUAN, J.-C., W. K. HÄRDLE, AND J. E. GENTLE, eds. (2012): Handbook of Computational Finance, Berlin, Heidelberg: Springer Berlin Heidelberg.

Risk-neutral density

Black & Scholes (1973)[3] showed that a continuously rebalanced stock-and-bond portfolio can replicate the payoff of a European option. This allows us to express the price of a European call as the discounted expected payoff under a risk-neutral probability density. Given call prices across strikes and a smooth interpolation, we can invert this relationship and recover the entire risk-neutral distribution implied by the market.

Let S_T be the underlying price at maturity T, r the risk-free rate, and $q(\cdot)$ the risk-neutral probability density. The starting price of a European call with strike K and maturity T is

$$C_0 = e^{-rT} \int_K^{\infty} (S_T - K)\, q(s)\, ds$$

where $q(s)$ denotes the risk-neutral density of S_T evaluated at the price level s.

Breeden & Litzenberger (1978)[4] showed that the second derivative of the call price with respect to the strike yields the RND:

$$e^{-rT} \frac{\partial^2 C(K,T)}{\partial K^2} = q(K)$$

This means that a smoothed call price function across strikes contains the market-implied probabilities of every future price level.

[3] BLACK, F. AND M. SCHOLES (1972): “The Valuation of Option Contracts and a Test of Market Efficiency,” The Journal of Finance, 27, 399. (1973): “The Pricing of Options and Corporate Liabilities,” The Journal of Political Economy, 81, 637– 654.

[4] BREEDEN, D. T. AND R. H. LITZENBERGER (1978): “Prices of State-Contingent Claims Implicit in Option Prices,” The Journal of Business, 51, 621.

Contact

GLOBAL PRODUCT LEAD FOR CREDIT
Davide Masi
T +44-20-78 62-72 67
davide.masi@eurex.com

VP FIC ETD PRODUCT
Leon von Essen
T +49-69-211-1 49 64
leon.von.essen@eurex.com

SALES EMEA
David Carretero
T +49-69-211-1 82 33
david.carretero@eurex.com

Lisa-Marie Zamfirescu
T +44-7484-06 34 55
lisa-marie.zamfirescu@eurex.com

SALES AMERICAS
Iris Hui
T +1-917-939-1624
iris.hui@eurex.com

SALES APAC
Neo Chiu
T +852-25 30-78 04
neo.chiu@eurex.com

Published by
Eurex Frankfurt AG
Mergenthalerallee 61
65760 Eschborn
Germany

www.eurex.com

ARBN Number
Eurex Frankfurt AG ARBN 100 999 764