MASTER

MASTER'S IN FINANCE

MASTER´S FINAL WORK

PROJECT

A CAPITAL-AT-RISK STRUCTURED PRODUCT ON GOLD AND DIGITAL-GOLD ASSETS

HAORAN SI

JUNE - 2026

MASTER

MASTER'S IN FINANCE

MASTER´S FINAL WORK

PROJECT

A CAPITAL-AT-RISK STRUCTURED PRODUCT ON GOLD AND DIGITAL-GOLD ASSETS

HAORAN SI

SUPERVISION:

JOÃO DUQUE

JUNE - 2026

Declaration on the Use of Generative AI

In preparing this work, the author made use of a generative AI tool in two areas. First, for the English-language translation and editing of text that was originally drafted by the author. Second, for assistance with the Python code, specifically: customising the Monte Carlo engine to the author's requirements, building the daily mark-to-market framework, and producing the final payoff output.

Acknowledgements

I would like to thank my supervisor, Professor João Duque, for his guidance, availability, and rigorous feedback throughout the development of this work. I am also grateful to the members of the jury for their constructive comments during the public examination. Finally, I thank my family for their unconditional support during my studies at ISEG.

Abstract

This work designs and assesses a three-year, multi-asset, capital-at-risk structured product aimed at retail investors. The product combines a fixed-income floor guaranteeing 70% of notional with an option-based return layer on four underlyings: a covered call on gold (GLD), a collar on a spot Bitcoin ETF (IBIT), a bull call spread on MicroStrategy (MSTR), and a ladder of barrier puts on the S&P 500 anchored to the market-wide circuit-breaker levels. Budgets are allocated by an equal-risk-contribution rule, which, given the near-zero realized correlation between gold and Bitcoin, reduces to inverse-volatility weighting. The economic role assumed for each underlying is tested against historical data before pricing; the "digital gold" hypothesis for Bitcoin is rejected and feeds back into the design. The product is priced by three independent methods — Black–Scholes closed forms, a binomial lattice, and Monte Carlo simulation — which agree at par, and its risk profile is characterised through vector and scalar Greeks. Real-world Monte Carlo scenarios, a historical backtest across three market regimes, and sensitivity analyses on the hedge weight, floor level, and volatility inputs show that the product's value lies in compressing the tail of the return distribution: it achieves a comparable Sharpe ratio at roughly one fifth of the volatility of naked holdings, at the cost of capped upside in bull markets. The product is not a better investment but a different risk–return trade-off — a clear-eyed exchange of upside for downside protection.

Keywords: Structured products; Capital at risk; Tail-risk hedging; Equal risk contribution; Monte Carlo simulation.

Resumo

Este trabalho concebe e avalia um produto estruturado multiativos a três anos, com capital em risco, dirigido a investidores de retalho. O produto combina um piso obrigacionista que garante 70% do nocional com uma camada de retorno baseada em opções sobre quatro ativos subjacentes: uma covered call sobre ouro (GLD), um collar sobre um ETF de Bitcoin à vista (IBIT), um bull call spread sobre a MicroStrategy (MSTR) e uma escada de opções de venda out-of-the-money sobre o S&P 500, ancorada nos níveis dos circuit breakers do mercado norte-americano. A alocação orçamental segue uma regra de contribuição igual de risco (ERC) que, dada a correlação realizada quase nula entre ouro e Bitcoin, se reduz a uma ponderação por volatilidade inversa. O papel económico assumido para cada subjacente é testado com dados históricos antes da avaliação; a hipótese do Bitcoin como "ouro digital" é rejeitada, com consequências diretas no desenho do produto. O produto é avaliado por três métodos independentes — fórmulas fechadas de Black–Scholes, uma árvore binomial e simulação de Monte Carlo — que convergem para o valor par, e o seu perfil de risco é caracterizado através de gregas vetoriais e escalares. Cenários de Monte Carlo, um backtest histórico em três regimes de mercado e análises de sensibilidade mostram que o valor do produto reside na compressão da cauda da distribuição de retornos: obtém um rácio de Sharpe comparável com cerca de um quinto da volatilidade das posições diretas, ao custo de um potencial de valorização limitado em mercados em alta. Não é um investimento melhor, mas um compromisso risco–retorno diferente.

Palavras-chave: Produtos estruturados; Capital em risco; Cobertura do risco de cauda; Contribuição igual de risco; Simulação de Monte Carlo.

Glossary

Abbreviation	Meaning
BS	Black–Scholes option pricing model (Black & Scholes, 1973)
CRR	Cox–Ross–Rubinstein binomial lattice (Cox et al, 1979)
CVaR	Conditional Value at Risk
DIP	Down-and-In Put
ERC	Equal Risk Contribution
ETF	Exchange-Traded Fund
GBM	Geometric Brownian Motion
GLD	SPDR Gold Shares ETF
IBIT	iShares Bitcoin Trust ETF
ITM / OTM	In / Out of The Money
MC	Monte Carlo (simulation)
MSTR	MicroStrategy Incorporated
MTM	Mark-to-Market
SPX	S&P 500 Index
VaR	Value at Risk
ZCB	Zero-Coupon Bond

Table of Contents

List of Figures

List of Tables

Chapter 1 – Introduction

1.1. Research Motivation and Research Question

Structured products hold an important place in the global retail investment market. Their value comes from embedding options, or combining different underlyings, to give investors a risk-return profile that sits between pure fixed income and pure equity. A traditional principal-protected note keeps the principal safe, but its participation in the upside is usually limited, and in a low-rate environment it is hard for such a note to produce an attractive return. At the same time, the rise of highly volatile assets such as cryptocurrency has created strong investor demand, but the sharp drawdowns of holding these assets directly — Bitcoin fell more than 65% in 2022 — keep many risk-averse investors away.

This tension raises a question that is both practical and of academic interest: can a multi-asset structured product be designed that offers meaningful upside participation while placing a hard limit on tail risk, and so does better, on a risk-adjusted basis, than a simple naked holding? And further: does such a product offer genuinely new value, or is it only a repackaging of instruments that already exist?

This work sets out to design and assess a multi-asset capital-at-risk structured product with a floor, aimed at retail investors, and to find a workable balance between these competing aims.

1.2. Product Overview

The product combines a floor guaranteeing 70% of notional with an option-based return layer on four underlyings — a covered call on GLD, a collar on IBIT, a bull call spread on MSTR, and a ladder of down-and-in puts on the S&P 500 — with budgets split by an equal-risk-contribution (ERC) rule. It is assumed issued on 30 January 2023 with a three-year term. Each underlying's assumed economic role is tested against historical data before pricing (Chapter 4); the full mechanics are set out in Chapter 3.

1.3. Contributions and Main Results

The contribution of this work is on two levels: methodological design and empirical testing.

Methodological contributions. This work (1) brings an equal-risk-contribution allocation into the design of a multi-asset structured product, which makes each return leg's risk budget transparent and measurable; (2) prices the product by three methods — Black–Scholes, a binomial lattice, and Monte Carlo — and checks that they agree; (3) in the Greeks analysis, separates the vector Greeks, one per underlying, from the scalar Greeks for the product as a whole, to identify the main sources of risk in a multi-asset setting; and (4) sets out a reusable way of studying a structured product, running from an ex-ante test of each asset's economic role, through pricing, risk analysis, multi-scenario simulation, and a multi-issue-date backtest, to a sensitivity check.

Main results. The three pricing methods agree at par and the daily mark-to-market path closes exactly on the maturity payoff. The product's value proves to lie in tail compression rather than in return: a comparable Sharpe ratio at a fraction of the volatility of naked holdings, bought by surrendering a large share of bull-market upside. Chapter 10 states the findings in full.

1.4. Structure of the Work

The work proceeds in four movements. Chapters 2 to 4 establish the ground: the literature the product draws on, its design, and an ex-ante test of the economic role assumed for each underlying. Chapters 5 and 6 value the product by three independent methods and measure its Greeks. Chapters 7 to 9 assess it — across Monte Carlo scenarios, against three historical issue dates, and under changes to its main design parameters. Chapters 10 and 11 conclude and set out the limitations.

Chapter 2 – Literature Review

This chapter reviews the research that the product draws on. It covers four areas. The first is the literature on structured products — how they are priced and whether they are good value for investors. The second is the debate on whether gold and Bitcoin act as safe havens, which is the basis for using both assets. The third is the work on tail-risk hedging, which motivates the down-and-in puts. The fourth is the work on multi-asset allocation, in particular risk parity, which is how the budget is split between the legs. The chapter ends by placing the product in this body of work.

2.1. Structured Products and Their Pricing

Structured products package a bond with one or more options to produce payoffs that plain instruments cannot. The familiar capital-protected note pairs a zero-coupon bond with a call, preserving the full principal; the product here protects only 70%, making it a capital-at-risk note that frees more budget for the option overlays. Whether such notes are fair value to the buyer is contested: Stoimenov & Wilkens (2005), examining the German market for equity-linked instruments, find issue prices systematically above theoretical value, which is one reason this work prices every leg independently rather than taking an issuer quote as given.

This work prices the product with three standard tools. The Black–Scholes model (Black & Scholes, 1973) gives closed-form values for the vanilla options. The binomial lattice of Cox et al (1979) provides a discrete-time check. A Monte Carlo simulation, following the methods set out in Glasserman (2004), gives a third, independent value and also handles the path-dependent parts of the payoff. For the barrier options, the closed-form solution of Reiner & Rubinstein (1991) is used, with the formulas as collected in Haug (2007). These methods are well established; the contribution here is not in the methods themselves but in applying them together to a multi-asset product and checking that they agree.

2.2. Gold and Bitcoin as Safe Havens

The product treats gold and Bitcoin as two different kinds of defensive asset, so the literature on their safe-haven properties is central. The starting point is the distinction drawn by Baur & Lucey (2010) between a hedge and a safe haven. A hedge is an asset that is uncorrelated with the market on average; a safe haven is an asset that holds its value specifically when the market falls sharply. The two are not the same, and an asset can be one without being the other. This distinction guides how the two assets are used in the product.

For gold the evidence is fairly settled: Baur & McDermott (2010), over three decades of data, find gold to be both a hedge and a safe haven for the major developed markets including the US, though not universally. Its safe-haven role is therefore real but conditional — strongest in developed markets and sharp downturns — which is why the product treats gold as a stable anchor whose protection is not assumed to hold in every condition.

For Bitcoin, the evidence is more divided, and the product is deliberately cautious about treating it as digital gold. Bouri et al (2017) find that Bitcoin is mainly a diversifier rather than a strong safe haven, and that even this role varies with the time horizon. Choi & Shin (2022) go further: Bitcoin acts as a hedge against inflation but not as a safe haven, and unlike gold its price falls when financial uncertainty rises, which is the opposite of safe-haven behaviour. Long et al (2021), using a nonlinear model, reach a similar conclusion, finding that gold can hedge several kinds of uncertainty while Bitcoin cannot. Chemkha et al (2021) study the COVID-19 period and find that gold was a weak safe haven during the crisis, while Bitcoin could not provide shelter because its own volatility rose too much. Taken together, this work points to Bitcoin being better seen as a high-volatility return asset than as a true safe haven. This is why the product does not rely on Bitcoin for protection — that role is left to the floor and to gold — but uses it for upside, with a collar to limit its large downside.

2.3. Tail-Risk Hedging

The tail hedge uses down-and-in puts. Ordinary diversification fails in a crisis, when correlations rise just as protection is most needed; a dedicated hedge paying off only in extreme down moves addresses this for a small premium. Bhansali & Davis (2010) argue that such a hedge, bought at the right price, need not be a pure drag — by capping the downside it can let an investor hold more risky assets, improving the whole portfolio's risk-return profile. The barriers are tied to the US circuit-breaker levels (−7%, −13%, −20%), regulator-set thresholds that mark widely recognised stress, so the hedge activates on events the market itself treats as severe. Because the puts are European, they target persistent downturns still depressed at maturity rather than brief crashes; this choice and its limitations are discussed in Chapter 11.

2.4. Multi-Asset Allocation and Risk Parity

With several option legs, a rule is needed to split the budget. Mean–variance optimisation relies on hard-to-estimate expected returns that make weights unstable; risk-based allocation avoids this. Maillard et al (2010) set out the equal-risk-contribution (ERC) portfolio, in which each component contributes an equal share of total risk; they show it is unique and that its volatility lies between the minimum-variance and equally weighted portfolios. Needing only a covariance matrix, ERC is a natural fit, and is used here to size the two core legs — gold and Bitcoin — so that neither dominates the sleeve's risk; the remaining legs are sized in the design chapter.

2.5. Positioning of This Work

Against this background, the product combines strands usually treated separately: the bond-plus-option structure of the structured-products literature, relaxed to a 70% floor to fund richer overlays; the safe-haven finding that gold and Bitcoin play different roles, each given its own leg; the equal-risk-contribution rule, sizing those legs by risk rather than by return forecasts; and a dedicated tail hedge tied to recognised stress levels. Whether the combination improves on the simple alternatives or merely repackages them is the question the rest of the work answers.

Chapter 3 – Product Design and Structure

This chapter describes the structured product that is priced and analyzed in the rest of the work. The product is a three-year, capital-at-risk note. It is issued on 30 January 2023 and matures on 30 January 2026. It combines a floor leg with a portfolio of options on four underlyings: the SPDR Gold Shares ETF (GLD), the iShares Bitcoin Trust ETF (IBIT), MicroStrategy Incorporated (MSTR), and the S&P 500 index (SPX). The aim of the design is a payoff that stays stable across bull, bear, and range-bound markets, while keeping the loss of principal limited.

3.1. Overall Architecture

The product is split into two parts: a floor leg and a risky leg. The floor leg secures a minimum redemption at maturity, which caps the downside. The risky leg uses a set of option overlays to add upside at maturity and to provide tail protection. If the notional is 100 and the issue-date cost of the floor leg is α, then the risky leg receives $100 - \alpha$. The minimum redemption is set at 70% of notional, and the floor is delivered by a zero-coupon bond, so its issue-date cost satisfies $\alpha = 70 \cdot e$^(−rT). Here r is the three-year US Treasury yield on the issue date and $T = 3$. This gives $\alpha = 62.19$, and $100 - \alpha = 37.81$. In other words, of every 100 units of notional, 62.19 units buy the zero-coupon bond and 37.81 units fund the risky legs. By accepting a loss of at most 30%, the product takes on a larger risky exposure than a fully bond-backed note would allow. This lets investors who want exposure to assets such as cryptocurrency, but are wary of their high volatility, take part while avoiding catastrophic losses.

3.2. Floor Leg

The floor leg is a three-year US zero-coupon bond that pays 70 units of notional at maturity. Its value on the issue date is the discounted redemption:

$$V_{\text{floor}} = 70 \cdot e^{-\text{rT}} = 70 \cdot e^{-0.039451 \cdot 3} = 62.19 \tag{3.1}$$

The floor leg has no optionality, so it contributes no delta, gamma, or vega to the product. Together with the tail-risk hedge described later, it forms the protective part of the structure.

A zero-coupon instrument is used rather than a coupon-paying bond so that the entire protective payment lands at maturity, matching the European settlement of every option leg and removing any reinvestment assumption. Two consequences follow. First, the cost of the guarantee is purely a function of the term rate: the 70-unit floor absorbs 62.19 of the notional at the 3.9451% issue-date rate, but 68.51 and 67.14 at the lower rates of the 2021 and 2020 issue dates (Table XII), so the floor is cheapest — and the product most generous — precisely when rates are high. Second, the 70 units are a credit claim on the bond issuer rather than a risk-free promise; the analysis here abstracts from issuer default.

3.3. Risky Leg

The risky leg is divided among four option-based strategies, one for each underlying. Each strategy is built so that its strikes are set as a percentage of the underlying's spot price on the issue date. This makes the design scale-invariant and reproducible at any issue date. Strikes are chosen mainly by computing the risk-neutral probability that the option ends in the money, $N(d_2)$ or $N(-d_2)$.

3.3.1. GLD Covered Call

The gold leg is long GLD with an out-of-the-money call sold at 118% of issue-date spot ($210.94); its unit value is $S_0 - C$. The premium raises the leg's effective return in exchange for capping gains above 118%. A protective put is deliberately omitted: the product-level floor already covers the downside, so a collar here would be redundant and costly.

3.3.2. IBIT Collar

The Bitcoin leg is a collar — long IBIT, a put bought at 60% of spot ($8.00), and a call sold at 150% ($20.00); the unit value is $S_0 + P - C$. Unlike the gold leg, this one keeps its put: on the most volatile underlying, leg-level protection remains valuable even with the floor in place, while the wide asymmetric band — a deep out-of-the-money put and a cap higher than GLD's 118% — preserves convex upside in risk-on phases.

3.3.3. MSTR Bull Call Spread

The MicroStrategy leg is a bull call spread — long the 110% call ($27.03), short the 300% call ($73.71); the unit value is $C(K_1) - C(K_2)$. It is a pure long-convexity overlay:

the maximum loss is the net premium paid, the payoff is capped at the spread width, and it adds leveraged Bitcoin convexity to the product at a small outlay.

3.3.4. SPX Layered DIP

The tail hedge is three layered down-and-in puts (DIPs) on the S&P 500. The barriers sit at 93%, 87%, and 80% of the issue-date index — the thresholds of the US market-wide circuit breakers (−7%, −13%, −20%) — so activation is tied to recognised levels of market stress; the strikes sit moderately out of the money at 90%, 85%, and 80%, keeping the hedge cheap in normal conditions.

Two design choices are worth noting. First, each barrier is set at or above its strike; a put that finishes in the money has then necessarily touched its barrier, so the knock-in condition is satisfied on every paying path and each tranche prices identically to a plain out-of-the-money put — the barrier anchors the design to the recognised circuit-breaker levels rather than adding a separate source of cheapness. Second, the puts are European and settle at maturity, so the hedge targets a persistent bear market — 2000–2002 or 2008 — rather than a flash crash that recovers before maturity; the limitations of this choice are discussed in Chapter 11.

3.4. Risk-Based Allocation

The correlation analysis shows that the realized correlation between Bitcoin and gold is weak, so the equal-risk-contribution (ERC) condition between them reduces to inverse-volatility weighting, with each weight proportional to the reciprocal of the asset's volatility. The equal-risk weights are computed using the Solver function in Excel. The DIP tail hedge and the MSTR option spread are assigned fixed weights of 10% and 5% of the risky sleeve, matching their roles as a dedicated hedge and a bounded convexity overlay and preventing an overlarge allocation to either. The equal-risk split is then applied to the remaining 85% core sleeve, between the gold and Bitcoin legs.

Applying this rule gives the final notional weights in Table I. The 10% DIP weight is itself the outcome of the comparison in the sensitivity analysis, which finds it preferable to higher allocations on a risk-adjusted basis.

TABLE I
FINAL NOTIONAL ALLOCATION AND CONVERSION FACTORS

Leg	Weight	Unit value	Conversion factor N
Floor (zero-coupon bond)	62.19%	62.19	—
GLD covered call	68.99%	163.53	0.159498
IBIT collar	16.01%	10.52	0.575750
MSTR bull call spread	5.00%	4.47	0.423080
SPX layered DIP	10.00%	781.44	0.004839

Source: Author's calculations. The weights for the four risky legs are shares of the risky sleeve. The floor leg (62.19) and the risky sleeve (37.81) sum to the 100-unit notional. The conversion factor N is the number of units of each leg, equal to its budget divided by its issue-date unit value.

It should be noted that the allocation has one limit. The ERC condition is imposed at the level of the underlying assets, using the volatilities of the underlyings rather than those of the option overlays. The option overlays compress the effective volatility of each leg asymmetrically — the IBIT collar truncates both the upside and the downside, while the GLD covered call truncates only the upside — so the actual risk contribution of the Bitcoin leg is lower than the underlying-level calculation suggests.

3.5. Payoff at Maturity

Putting the parts together, the redemption value of the product at maturity is the sum of the floor redemption and the value of each risky leg, scaled by its conversion factor. Let S^T denote the maturity prices of the four underlyings and m the lowest level reached by the S&P over the life of the note. The per-unit leg payoffs are:

GLD: $\min(S_T^{GLD}, K_c)$ (3.2)

IBIT: $\mathrm{clip}(S_T^{IBIT}, K_p, K_c)$ (3.3)

MSTR: $\max(0, \min(S_T^{MsTR}, K_s) - K_l)$ (3.4)

DIP: $\Sigma_j \, 1\{m \leq B_j\} \cdot \max(K_j - S_T^{spx}, 0)$ (3.5)

where the indicator $1\{\cdot\}$ captures the knock-in condition of each DIP tranche j. The total payoff at maturity is then:

$$V_T = 70 + N_{GLD} \cdot GLD + N_{IBIT} \cdot IBIT + N_{MSTR} \cdot MSTR + N_{DIP} \cdot DIP \quad (3.6)$$

The floor guarantees a minimum redemption of 70 units, whatever the risky legs do, and the layered DIP adds to the floor in a persistent downturn. The short calls in the gold, Bitcoin, and MicroStrategy legs cap the product's participation in a strong rally. This is the central trade-off of the structure: it gives up extreme upside in exchange for downside protection and option carry. The pricing of each leg, and the validation of the resulting product value, are the subject of Chapter 5.

Chapter 4 – Underlying Economic Roles and Hypothesis Testing

Having set out the product design in the previous chapter, this chapter checks that the choice of the four underlying assets is sound, and explains the logic behind selecting GLD, IBIT, MSTR, and SPX. This is an ex-ante check, carried out before pricing: each underlying is given an assumed economic role, and that assumption is then tested against historical data. Three hypotheses are examined — IBIT as digital gold, MSTR as a leveraged proxy for Bitcoin, and the stability of gold's safe-haven property. As will be seen, one of them is rejected by the data, and the test results feed back into the product design.

4.1. Hypothesis 1: Is IBIT a Digital Gold? — Rejected

The market narrative treats Bitcoin as 'digital gold' — an inflation hedge and quasisafe-haven. If so, IBIT should correlate positively with GLD. Over 2023–2026 the correlation of daily returns is ρ(GLD, IBIT) = 0.026: the hypothesis is rejected. IBIT moves essentially independently of gold and behaves as a high-volatility risk asset tied to macro liquidity, in line with Choi & Shin (2022). The design consequence is direct: IBIT is treated as a high-beta return leg held inside a collar, not as a defensive asset — the safe-haven role stays with the floor and the GLD leg.

4.2. Hypothesis 2: MSTR as a Leveraged Proxy for Bitcoin — Confirmed but Time-Varying

Since its August 2020 treasury pivot, MicroStrategy's balance sheet is structurally long Bitcoin, so MSTR should act as a leveraged proxy for IBIT. The sensitivity β of MSTR returns to Bitcoin returns is about −0.01 before the pivot and about 0.52 after it: the hypothesis holds for the product's issue period, but the link is time-varying. The leg is therefore kept small, paired with IBIT on the offensive side, and its Monte Carlo drift is derived as a multiple of the IBIT drift.

4.3. Hypothesis 3: Is Gold's Safe-Haven Property Stable? — Confirmed but Regime-Dependent

Gold's safe-haven property is widely accepted; the sharper question is whether its low correlation with equities survives a crisis. Splitting the sample into normal and crisis windows (90 days around the 2020 flash crash) gives ρ(GLD, SPX) = 0.02 in calm markets but 0.23 in the crisis: the property is regime-dependent, in line with Baur & McDermott (2010). The product therefore never relies on GLD alone for protection — the fixed-income floor and the DIP ladder provide the baseline guarantee even when gold's safe-haven purity fades.

Taken together, one hypothesis was rejected and two confirmed with qualifications: the underlyings enter the product with economic roles set out beforehand and tested against data, and the rejection itself improved the design.

Chapter 5 – Pricing Methodology

Each part of the product is priced separately by three cross-checked methods — Black–Scholes, a binomial lattice, and Monte Carlo — with the down-and-in puts priced on the Reiner & Rubinstein (1991) formula. All legs are valued at the issue date, 30 January 2023, over a three-year horizon at a risk-free rate of $r = 3.9451\%$.

5.1. Black–Scholes Pricing

The vanilla legs are priced with the Black & Scholes (1973) model. For an underlying with spot S, strike K, volatility σ, maturity T, and risk-free rate r, the price of a European call is:

$$C = S\,N(d_1) - Ke^{-rT}\,N(d_2) \tag{5.1}$$

with

$$d_1 = [\ln(S/K) + (r + \sigma^2/2)T] \,/\, (\sigma\sqrt{T}), \quad d_2 = d_1 - \sigma\sqrt{T} \tag{5.2}$$

and the put obtained by put–call parity. The four underlyings are priced with the issue-date parameters in Table II.

TABLE II
ISSUE-DATE PRICING PARAMETERS

Parameter	**GLD**	**IBIT**	**MSTR**	**SPX**
Spot S_0	178.76	13.3348	24.569	4,017.77
Volatility σ	15.26%	65.71%	110.82%	24.14%
Dividend yield q	0	0	0	1.70%

Source: Volatilities are trailing one-year realized volatilities, estimated from daily log returns to the issue date. The risk-free rate is the three-year US Treasury yield of 3.9451%.

One value needs a note. IBIT did not begin trading until January 2024, so the 65.71% volatility is derived from a synthetic series reconstructed from Bitcoin spot — a proxy, not a directly observed market value. This is the concession required by the no-look-ahead principle, under which only information available at issuance is used.

5.1.1. Vanilla Legs

Each vanilla leg is the underlying plus one or two European options: the gold covered call as spot minus the short call; the IBIT collar as spot plus the long put minus the short call; the MSTR bull spread as the long call minus the short call. The floor leg, at 62.19, is the largest single component of the product value and the main source of interest-rate risk.

5.1.2. Down-and-In Puts

The tail-risk hedge is made up of three layered down-and-in put tranches on the S&P 500. A down-and-in put pays the same as an ordinary put at maturity, but only if the underlying has touched a lower barrier B at some point during the option's life; otherwise it expires worthless. Since each tranche's barrier is set at or above its strike, the closed-form value follows Reiner & Rubinstein (1991). Defining

$$\lambda = (r - q + \sigma^2/2) / \sigma^2 \tag{5.3}$$

and the auxiliary terms x_1, y, and y_1 as in Haug (2007), the value of a down-and-in put with the barrier at or above the strike is:

$$\begin{aligned} p_{\mathrm{di}} = {} & -Se^{-qT}N(-x_1) + Ke^{-rT}N(-x_1+\sigma\sqrt{T}) \\ & + Se^{-qT}(B/S)^{2\lambda}\,[N(y) - N(y_1)] - Ke^{-rT}(B/S)^{2\lambda-2}\,[N(y-\sigma\sqrt{T}) - N(y_1-\sigma\sqrt{T})] \end{aligned} \tag{5.4}$$

The auxiliary terms are defined in Appendix C. The three tranches have barriers at

93%, 87%, and 80% of the issue-date index level, and strikes at 90%, 85%, and 80%, matching the thresholds of the US circuit breakers. Pricing them with the S&P parameters gives tranche values of 327.48, 257.33, and 196.62 index points, for a combined unit value of 781.44.

Because the puts are monitored discretely (daily) rather than continuously, the continuous-barrier formula in principle overstates the knock-in probability. Applying the Broadie et al (1997) continuity correction shifts each barrier by a factor in the volatility and the monitoring frequency. For daily monitoring over a three-year horizon, the resulting change in value is below 0.01% and is neglected. The closed-form tranche values were also cross-checked against a Monte Carlo estimate with daily barrier monitoring, with agreement to within simulation error.

5.2. Value Additivity

Summing the floor and the four legs, each scaled by its conversion factor, gives a product value of 100.00 — par. This equality is an identity, not a validation: each conversion factor is defined as the leg's budget divided by its unit value, so the sum returns the notional by construction. Table III reports the resulting leg values. It confirms only that the allocation was implemented without arithmetic error; the substantive check of the option values themselves lies in the agreement across the three independent methods below, and in the coherence of the Greeks in the next chapter.

TABLE III
LEG VALUES AT ISSUANCE (PER 100 NOTIONAL)

Leg	**Unit value**	**Value (N × unit)**
Floor (zero-coupon bond)	62.19	62.19
GLD covered call	163.53	26.08
IBIT collar	10.52	6.06

MSTR bull call spread	4.47	1.89
SPX layered DIP	781.44	3.78
Product	—	**100.00**

Source: Author's calculations. Unit values are issue-date Black–Scholes prices; the product value is the budget-weighted sum and equals par by construction.

5.3. Binomial Validation

Each option is re-priced on a Cox et al (1979) binomial lattice as an independent check. Over T split into n steps of length $\Delta t = T/n$, the underlying moves up by $u = \exp(\sigma\sqrt{\Delta t})$ or down by $d = 1/u$, with risk-neutral up-probability $p = (e^{(r-q)\Delta t} - d)/(u - d)$ (5.5). All options are European, so the lattice is rolled back by simple discounting, with no early-exercise test. At n = 36 the five option prices sit within 0.75% of their closed-form values, with the oscillatory convergence typical of the method; at the product level the lattice gives 100.01 against the closed-form 100.00 — a one-basis-point model difference, not a numerical error.

5.4. Monte Carlo Validation

The third method simulates the joint evolution of the four underlyings and prices the product by averaging the discounted payoffs. Correlated paths are generated by a Cholesky decomposition of the correlation matrix, which preserves the correlation among the underlyings. Under the risk-neutral measure, each underlying follows geometric Brownian motion with drift $r - q$, simulated on a daily grid over the three-year horizon. The minimum of the simulated S&P path is tracked along each trajectory to determine knock-in of the down-and-in puts.

With 100,000 paths, averaging the payoffs prices the product at 100.0061, close to the closed-form and binomial values. In particular, the simulation reproduces the down-and-in puts through explicit daily barrier monitoring, which provides an independent check on the Reiner–Rubinstein formula without relying on its closed-form assumptions. Table IV collects the three estimates.

TABLE IV
PRODUCT VALUE BY METHOD

Method	**Product value**
Black–Scholes (closed form)	100.00
Binomial lattice (36 steps)	100.01
Monte Carlo (100,000 paths)	100.0061

Source: Author's calculations. All three methods price the note within two basis points of par.

The three methods agree to within two basis points — the central pricing result. Their roles are complementary: Black–Scholes gives closed-form values, the lattice an independent discrete-time check, and Monte Carlo both a third value and the path-dependent barrier logic.

Chapter 6 – Greeks and Risk Sensitivities

The vanilla legs have closed-form Greeks; those of the down-and-in puts are obtained by bump-and-revalue. The sensitivities fall into two kinds: those defined against a single underlying, and those shared across all legs.

6.1. Closed-Form Greeks for the Vanilla Legs

For the single European options in the gold, Bitcoin, and MicroStrategy legs, the Greeks follow from the Black–Scholes formula (Black & Scholes, 1973). Writing φ for the standard normal density and N for its cumulative distribution, the call Greeks are delta $= N(d_1)$, gamma $= \varphi(d_1)/(S\sigma\sqrt{T})$, vega $= S\varphi(d_1)\sqrt{T}$, together with theta and rho, and the put values follow by parity. The Greeks of each leg are then built up from the option Greeks according to how the leg is put together. The covered call is long the underlying and short a call, so its delta is $1 - N(d_1)$, and it takes on the short call's negative gamma and negative vega. The collar combines the underlying with a long put and a short call, and the bull call spread nets the Greeks of its two calls.

Computing these at the issue-date parameters gives the leg-level sensitivities in Table V. The signs read directly from the short calls: each adds negative gamma and negative vega, leaving the three return legs short volatility; the SPX hedge leg pulls the other way, and the net position is set out in Section 6.4. In the table, delta is sensitivity to the underlying price, gamma the rate of change of delta, vega the sensitivity to volatility, and theta to the passage of time (per year).

TABLE V
LEG-LEVEL GREEKS AT ISSUANCE

Leg	**Delta**	**Gamma**	**Vega**	**Theta**
GLD covered call	0.5185	−0.00844	−123.39	+5.93
IBIT collar	0.2447	−0.011	−3.85	+0.708
MSTR spread	0.1611	−0.00237	−4.76	+0.899

Source: Author's calculations. Values are per unit of the respective leg, expressed in the leg's own underlying. Vega is per +1.00 (100 percentage points) of volatility; theta is per year. Delta, gamma, and vega are not additive across legs (see Section 6.3).

6.2. Greeks of the Down-and-In Puts by Bump-and-Revalue

The Greeks of the down-and-in puts are hard to obtain in closed form, so they are computed by a bump-and-revalue method. Each input that a Greek measures is changed by a small amount, the tranche is re-priced, and the Greek is read off as the resulting difference. A central difference is used for delta and gamma, taking one bump up and one bump down for second-order accuracy.

One detail of the bump is essential. The barrier and the strike of each tranche are both defined relative to the issue-date S&P level, denoted *S*ref, and are held fixed during the bump; only the current price, *S*val, is changed. Bumping the reference level along with the spot would move the barrier itself and corrupt the sensitivity, so the two are kept separate. The auxiliary exponent λ of the Reiner–Rubinstein formula is a function of the spot and is not bumped on its own; it is recomputed automatically when the tranche is re-priced at the changed spot. With this separation observed, the down-and-in put delta is small and negative at issuance, consistent with a deep-out-of-the-money put that gains value as the index falls toward its barriers.

6.3. Vector and Scalar Greeks

A multi-asset product has sensitivities of two kinds. Delta, gamma, and vega are each defined against a particular underlying — a gold delta and a Bitcoin delta lie along different dimensions and cannot be summed into one product-level number, so they are reported as a vector, one entry per underlying (Table VI). Theta and rho are different: every leg shares the same calendar and, through discounting, the same risk-free rate, so these sensitivities are comparable across legs and add into a single scalar. This is why Table VI reports product theta and rho as totals while the deltas, gammas, and vegas are not summed.

TABLE VI
PRODUCT GREEKS: VECTOR COMPONENTS AND SCALAR TOTALS

Vector Greek	**GLD**	**IBIT**	**MSTR**	**SPX**
Delta	+0.08269	+0.14090	+0.06816	−0.00310
Vega	−0.1968	−0.0222	−0.0201	+0.2857
Scalar Greek (product total)		**Value**		
Theta (per year)		+3.47		
Rho (per +1 percentage point)		−2.82		

Source: Author's calculations. Vector Greeks are reported per underlying and are not summed across columns; scalar Greeks are product-level totals. Vega here is per +1 percentage point of volatility, matching the units of rho. Gamma components are omitted from the table for brevity.

6.4. The Product's Risk Profile

The product's overall risk profile has three notable features. First, it earns positive carry: its annualised theta is +3.47, because the pull-to-par of the floor and the time decay collected on the short calls together outweigh the decay paid on the down-and-in puts. All else equal, the note gains value as time passes — an attractive property for a protective structure.

Second, the product's volatility exposure is two-sided: the three return legs are short gamma and short volatility through their embedded short calls, while the SPX hedge leg is long volatility. Which effect dominates depends on which volatility moves — and in a crisis it is typically the S&P volatility that rises most, so the hedge leg gains precisely when it matters. This is the natural counterpart of the carry: in effect, the note is compensated through theta for bearing short-volatility risk. Third, among the scalar sensitivities, rho is the dominant market-risk exposure, which reflects the floor leg's role as the main source of interest-rate duration. Among the vector sensitivities, the S&P enters most strongly, through the down-and-in puts and, indirectly, through its high correlation with MicroStrategy. The hedging implication is that interest-rate risk can be managed at the product level as a single exposure, while the equity and crypto deltas must be hedged underlying by underlying.

Chapter 7 – Monte Carlo Investment Analysis

Building on the pricing completed in Chapter 5, this chapter replaces the risk-neutral measure used there and simulates the product under real-world return assumptions, describing its return distribution and comparing it with naked holdings to assess its downside protection. The approach is the same as in Section 5.4.

7.1. Simulation Framework

The difference from Section 5.4 is that the drift is changed from the risk-free rate *r* to each asset's real-world expected return. The expected returns of GLD, IBIT, and SPX are reasonable assumptions set according to each market scenario. Since MicroStrategy became a leveraged proxy for IBIT after it adopted its Bitcoin treasury strategy in August 2020, its drift is set to 1.5 times the IBIT drift plus 0.02, where the 0.02 reflects its remaining software business. The one exception is the risk-neutral scenario, in which every drift is set to the risk-free rate. Each scenario is simulated over 100,000 daily paths with a fixed seed.

7.2. Scenarios and Results

The market scenarios and their drifts are shown in Table VII. They cover a benign environment, a crypto-led environment, a bear-market and recession environment, and a stagflationary environment.

TABLE VII
SCENARIO DRIFT ASSUMPTIONS (ANNUALISED EXPECTED RETURNS)

Scenario	**GLD**	**IBIT**	**MSTR**	**SPX**
Base	6%	15%	24.5%	9%
CryptoBull	3%	60%	92%	12%
BearCrisis	8%	−40%	−58%	−15%
Stagflation	15%	5%	9.5%	−2%
RiskNeutral	3.95%	3.95%	3.95%	3.95%

Source: Author's assumptions. The MSTR drift is derived as 1.5 times the IBIT drift plus 0.02. In the risk-neutral scenario every drift equals the risk-free rate.

The return distributions across the five scenarios are reported in Table VIII. In every assumed market environment, the annualised return is higher than in the risk-neutral environment, while the volatility stays inside a narrow band of roughly 4.6% to 5.2%, close to the risk-neutral level in every case. This is the result of the floor leg and the double-sided cap of the collar, which give the product a stable holding profile. The probability of loss is also very low: because of the protective structure, the loss probability stays low even in extreme environments, while in the risk-neutral environment the decay of option time value leads to a higher loss probability. One result deserves comment: the expected return is highest in the BearCrisis scenario (6.98%), above both the Base (4.35%) and the CryptoBull (6.71%) cases. This is not an artefact. The floor removes the downside that a falling market would otherwise inflict, while the down-and-in puts knock in and pay, so a severe fall converts into a positive contribution rather than a loss. The structure is therefore convex to the downside by construction, and the scenario ranking should be read as evidence that the tail hedge works, not that the product is a directional bear-market bet: the same design gives up most of the upside in the CryptoBull case, where the short calls cap all three return legs.

TABLE VIII
PRODUCT RETURN DISTRIBUTION BY SCENARIO

Scenario	**E[ann]%**	**Vol%**	**Sharpe**	**VaR95%**	**CVaR95%**	**P(loss)%**	**Median**
Base	4.35	4.80	0.084	−0.09	−2.90	5.1	112.15
CryptoBull	6.71	5.13	0.539	3.48	0.07	2.0	123.27
BearCrisis	6.98	5.16	0.587	6.56	2.75	0.9	123.40
Stagflation	6.04	4.66	0.449	8.25	7.32	0.3	117.64
RiskNeutral	3.95	4.84	0.000	−1.49	−4.16	7.4	110.62

Source: Author's calculations, 100,000 simulated paths per scenario. E[ann] and Vol are annualised; P(loss) is the probability that the terminal value falls below the 100 notional.

The terminal-value distributions are shown in Figure 3 (Appendix C). Each distribution is concentrated above par, with a clear spike at the cap.

7.3. Comparison with Naked Holdings (Base Scenario)

To judge whether the structured product really creates value, it is compared in the base scenario with three naked alternatives: holding GLD only, holding SPX only, and an equal-weight holding of the four underlyings (25% each). Table IX shows the comparison.

TABLE IX
PRODUCT VERSUS NAKED HOLDINGS (BASE SCENARIO)

Strategy	**E[ann]%**	**Vol%**	**Sharpe**	**CVaR95%**	**Worst**
Structured product	4.35	4.80	0.084	−2.90	85.1
GLD only	5.39	15.29	0.094	−32.85	35.0
SPX only	5.47	24.17	0.063	−51.33	20.2
Equal-weight	6.91	36.20	0.082	−54.03	24.1

Source: Author's calculations, 100,000 simulated paths. Worst is the minimum terminal value across paths, on an invested base of 100.

The four Sharpe ratios are close, but the product earns 4.35% against 5.4% to 6.9% for the naked alternatives while running about one fifth of their volatility, and its conditional value at risk (−2.9) is far smaller than the naked alternatives (−33 to −54). This contrast is shown in Figure 1: the product's payoff distribution is far narrower, its mass concentrated just above the notional, while the two naked positions spread much wider. Figure 4 (Appendix C) shows the product's payoff against naked exposure: it holds up when the market falls sharply, while the naked positions drop throughout, at the cost of a capped upside.

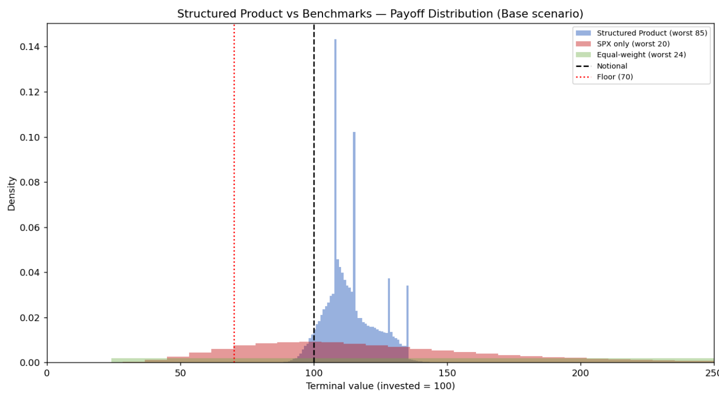

FIGURE 1 – Structured product versus naked holdings, payoff distribution (base scenario).

7.4. Downside Protection under Extreme Stress

The product's value is also tested directly under an extreme bear-market stress scenario, in which Bitcoin and the US equity market fall sharply at the same time. Table X shows the product against a naked S&P 500 position and an equal-weight holding of the four underlyings.

TABLE X
DOWNSIDE PROTECTION (BEAR-CRISIS SCENARIO)

Strategy	**Median**	**Worst**	**P(below 70)**	**CVaR95%**
Structured product	123.4	87.2	0.00%	2.7
SPX only	55.5	9.8	—	—
Equal-weight	53.6	17.4	—	—

Source: Author's calculations, 100,000 simulated paths. P(below 70) is the probability the terminal value falls below the floor.

The product's median terminal value is 123, its worst simulated outcome is 87, and the probability of breaching the 70 floor is zero. By comparison, the worst outcome of the naked S&P 500 position is close to 10, and that of the equal-weight portfolio is 17. The floor and the layered down-and-in puts together hold the product's loss far above the level of direct investment.

Taking the simulation results together, the product's value proposition is established. It is not designed to beat naked exposure in a rising market; rather, it works by compressing the return distribution — keeping an expected return above the risk-free rate while avoiding the large losses that direct investment in these volatile underlyings would bring.

Chapter 8 – Historical Backtest

This chapter shows the product's performance in practice, from issuance to maturity. The 2023–2026 period is already history, and it was a bull market that was very favourable to both crypto and safe-haven assets. Using this window as a validation would make any backtest look good and would lose its meaning as a test. The main argument was made in the previous chapter; this chapter is positioned as a supplement that shows the product along a single path that really occurred.

8.1. Single-Issue-Date Backtest (Issued 30/01/2023)

We take the product issued on 30 January 2023 as the base case, and simulate the full process of holding it to maturity on 30 January 2026. All pricing parameters — volatility, the risk-free rate, strikes, barriers, and position sizes — are computed from the trailing data available on the issue date, with no future information introduced.

8.1.1. Settlement at Maturity

At maturity all four underlyings had risen (Table XVI, Appendix C), and the strong bull market pushed all three return legs to their caps. The S&P's lowest point over the period was 3,855.76 (13/03/2023), only 96% of the issue-date level, so the DIP barrier was never touched and the tail hedge expired worthless — the cost of protection that, in a bull market, is spent in vain.

The product payoff at maturity is therefore set out in Table XI. Applying the payoff formula of Chapter 3, with each leg payoff scaled by its conversion factor, gives:

TABLE XI
PRODUCT PAYOFF AT MATURITY (SINGLE ISSUE DATE)

Leg	**Payoff**	**Note**
GLD covered call	210.94	capped
IBIT collar	20.00	capped
MSTR spread	46.68	max
DIP	0	not triggered
Floor	70	
Product payoff	**134.91**	**annual +10.50%**

Source: Author's calculations. The product payoff is the floor plus each leg payoff scaled by its conversion factor.

$$V_T = 70 + 0.159498 \times 210.94 + 0.575750 \times 20 + 0.423080 \times 46.68 + 0.004839 \times 0 \quad (8.1)$$

$$= 70 + 33.64 + 11.52 + 19.75 + 0 = 134.91 \tag{8.2}$$

The payoff is 134.91, an annualised 10.50% (compared with naked holdings over the same period in Table XVII, Appendix C).

As expected, the product underperformed every naked holding in this bull market: the short-call caps gave up most of the upside. This is the structure working as designed — its value is downside protection, not directional return — and matches its positioning for conservative investors seeking measured crypto exposure.

8.1.2. Daily Mark-to-Market

To capture the swings during the holding period, a daily mark-to-market simulation re-prices the product on each trading day, holding the pricing parameters fixed and varying only the spot levels and remaining time.

The result is shown in Figure 2. The values at the issue date and the maturity date cross-check against the earlier payoff, and the volatility is very low. The maximum drawdown (1.9%) occurs only at the very start of the product's life, which further shows that, even though the final return is below naked holdings, the holding process is almost painless.

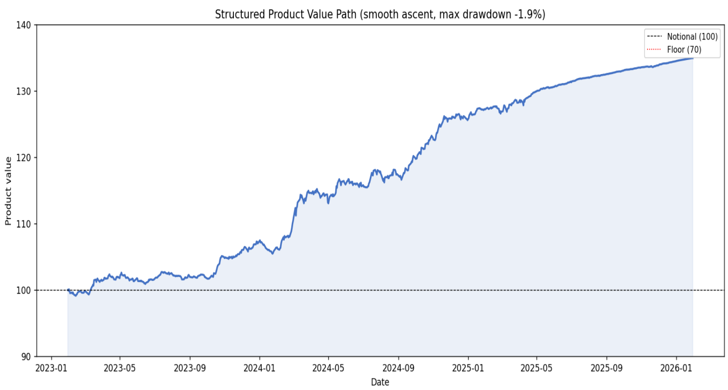

FIGURE 2 – Structured product value path (daily mark-to-market).

Figure 5 (Appendix C) shows the product against naked holdings, where the product's low-volatility feature is clear.

8.2. Robustness across Multiple Issue Dates

The same product is issued under the same rules at two further historical points — 08/11/2021 (the crypto peak, followed by several years of bear market) and 19/02/2020 (the COVID-19 flash crash and circuit breaker in US equities) — with S_0, σ, r, strikes, barriers, and N all recomputed from the conditions prevailing on each date. Table XII summarises the results.

TABLE XII
BACKTEST ACROSS THREE ISSUE DATES

	2023-01-30	**2021-11-08**	**2020-02-19**
Environment	Bull base	Crypto peak	Pre-COVID
Maturity	2026-01-30	2024-11-08	2023-02-19
r	3.95%	0.72%	1.39%
alpha (floor)	0.6219	0.6851	0.6714
S&P trough	96.0%	76.1%	66.1%
DIP	not triggered	knocked in	knocked in
Product payoff	134.91	119.33	114.34
Annualised	+10.50%	+6.07%	+4.57%
Max drawdown	−1.9%	−7.2%	−15.6%
Equal-weight	347	178	165

Source: Author's calculations. All parameters are recomputed from trailing data on each issue date. The equal-weight benchmark is the terminal value of a 25% holding of each underlying.

Across three different market cycles the product returned positively in every case, so its earlier result does not depend on a bull-market start. On both crisis issue dates the DIP barrier was touched, showing that the protection works outside simulation. In both cycles the product again underperformed the four naked holdings — the price of that protection — yet even in the COVID-19 flash crash its maximum drawdown stayed at 15.6%, without touching the floor.

Note too that the lower rates on the other two issue dates make the floor leg consume more of the budget when discounted to the start, further limiting upside — the natural effect of interest rates on the design.

Figure 6 (Appendix C) shows the product's daily mark-to-market payoff across the three historical periods, making the path more visual.

8.3. Discussion of Look-Ahead Bias

Look-ahead bias must be faced directly: GLD, IBIT, and MSTR were all ex-post winners over 2023–2026, so a long strategy naturally backtests well — a real risk, and I acknowledge it.

But neither the design logic nor the conclusions of the backtest depend on picking future winners, for two reasons.

First, the product underperforms naked holdings in a bull market. Had the aim been to look good in a backtest, the short-call caps, the floor, and the DIP would never have been built in, since all three are a drag when markets rise. Second, the underlyings were fixed in advance for their economic roles — gold as safe haven, IBIT as crypto exposure, MSTR as a leveraged proxy, SPX as the broad market — and for diversification, not for their realised returns.

Chapter 9 – Sensitivity Analysis

What is still missing is a check on the design itself: are the main parameters set in a sensible range, and how does the product's risk–return profile move if the DIP weight, the floor level, or the volatility of one of the underlyings changes?

This chapter answers these questions along three dimensions.

9.1. DIP Weight Sensitivity (10% / 20% / 30%)

The DIP weight is set in turn to 10%, 20%, and 30% of the risky sleeve. The MSTR spread is held at 5%, and the gold and Bitcoin legs split the remaining budget by the same equal-risk rule used earlier. Each case is re-simulated in the Monte Carlo framework of Chapter 7, and the expected return and Sharpe ratio are compared across the three weights. The results are in Table XIII (Appendix C).

The table shows a case of over-hedging. Raising the DIP weight from 10% to 20% and then 30% does not improve the product's risk-adjusted return; instead the Sharpe ratio falls, from +0.084 to −0.042 and then to −0.123. The annualised return drops as the weight rises. The BearCrisis column shows why this is not a fair trade: additional DIP does raise the expected return in a crisis (from 6.98% to 13.03%), but it does not improve the outcome that matters for a protective structure — the worst simulated terminal value actually falls, from 87.2 to 81.5, because the extra premium is paid on every path while the barrier is touched on only a few.

This looks odd at first. Why would more tail protection lower the product's performance? The reason is that the floor already does most of the downside work. The hard floor at 70% of notional already holds the product up in a severe bear market. On top of that, each extra unit of DIP buys little further protection, but its premium is still paid in full. A higher DIP weight therefore keeps eating into the return without cutting risk by a matching amount.

9.2. Floor-Level Sensitivity (50% / 70% / 90%)

Holding the other parameters fixed, the floor is set in turn to 50%, 70%, and 90% of notional. The aim is to see how the size of the risky sleeve, and the product's performance, change with the floor. Table XIV (Appendix C) reports the budget split in each case.

The floor and the risky sleeve trade off one for one. Raising the floor to 90% costs 79.95 of the notional to secure at the issue-date rate, leaving a risky sleeve of only 20.05: the GLD budget falls from 26.08 to 13.83 and the DIP budget from 3.78 to 2.00, so both the upside participation and the tail hedge are cut by roughly half. Lowering the floor to 50% costs 44.42 and frees a sleeve of 55.58, raising every leg budget by about 47%.

What the two sides buy is not symmetric: the guaranteed minimum moves linearly with the floor — 50, 70, or 90 units returned in the worst case — while the option budget that generates the upside moves the other way and far faster in proportional terms. The 70% level is retained because it is where the sleeve still funds all four legs at a meaningful size while the guaranteed minimum caps the loss at 30% of notional.

9.3. IBIT Volatility Sensitivity: Real (~50%) vs Synthetic (65.71%)

IBIT did not begin trading until January 2024. For the issue date in early 2023, and for the dates before it, the IBIT volatility used in pricing and in the backtest comes from a synthetic series reconstructed from Bitcoin. That synthetic volatility is 65.71%, while the volatility observed after IBIT started trading is about 50%. This section measures how much this difference matters for the pricing of the product.

Holding the other parameters fixed, the IBIT volatility is changed from 65.71% to 50%, and the collar's unit value and the matching number of units N are recomputed. Table XV (Appendix C) shows the two cases.

The higher synthetic volatility (65.71%) prices the collar lower, at 10.52 against 11.17. With the same budget, this lets the product hold more IBIT, so N is larger. Switching to the real volatility (50%) raises the collar's value, the number of units falls, and the product's IBIT upside is a little lower. The effect on the final result is small: as

Table XV (Appendix C) shows, the collar is about 6% dearer at the real volatility, and the IBIT position about 6% smaller.

Chapter 10 – Conclusion

10.1. Summary of the Research

This work designed a retail-oriented, capital-at-risk structured product — a 70% fixed-income floor plus four option-based return legs — and assessed it end to end. Three pricing methods agree at par; the Greeks show a product short volatility on its return legs and long volatility in its SPX hedge, earning positive carry, with interest rate as its dominant scalar exposure. The ex-ante tests rejected the 'digital gold' view of IBIT, and that rejection fed straight back into the design, placing IBIT inside a collar. The sensitivity tests confirmed the 10% DIP weight and the 70% floor as the balance points, and quantified the synthetic-IBIT volatility gap (65.71% against about 50% observed after listing).

The Monte Carlo analysis and the three-regime backtest together fixed the product's position: Sharpe ratios in the same range as the naked alternatives at roughly one fifth of their volatility, and a conditional value at risk of −2.9 against −33 to −54, paid for with bull-market upside.

10.2. Answering the Research Question

Does the product offer genuine value, or is it only a repackaging of existing instruments? It offers something a naked holding cannot. The floor and the down-and-in puts delivered real downside protection — in the 2021 and 2020 backtest crises and in the Monte Carlo stress scenarios — turning catastrophic losses into controlled drawdowns, with a largest backtest drawdown of only −1.9% to −15.6%. The cost is upside: in the

2023–2026 bull market the terminal value of 134.91 fell far below every naked strategy (173 to 609). Across the two dimensions of risk-adjusted return and tail-risk control the product does offer what a naked holding does not, but the value depends on risk preference — a clear-eyed trade-off, not a free lunch. The contribution is to combine known option structures and asset classes into one coherent product and put it through a full treatment, from ex-ante hypothesis testing to sensitivity checks: a framework that may serve both academic work on structured products and their practical design.

Chapter 11 – Limitations and Future Outlook

This chapter sets out the limitations of the work on three levels — data, method, and scope — together with the direction of each effect and how it could be addressed in future work.

11.1. Data Limitations

The synthetic IBIT series. IBIT did not list until January 2024, while the backtests reach back to 2023, 2021, and 2020; the synthetic series built from the Bitcoin spot price bridges this gap but carries assumption bias. It implies a volatility of 65.71%, against roughly 50% observed after listing; Chapter 9 quantifies the consequence — at the real volatility the collar's unit value rises from 10.52 to about 11.17, lowering the IBIT position by about 6% for the same budget, so the product's upside participation may be overstated. The proxy is weakest for the earliest years, when the Bitcoin market lacked institutional participation and behaved fundamentally differently from 2023–2024. Once IBIT has a longer real history, the pricing, backtest, and sensitivity analysis can be redone on fully real data.

MSTR's role in the 2020 backtest. The multi-issue-date backtest includes an issue date of February 2020, but MicroStrategy adopted its Bitcoin treasury strategy only in

August 2020; at that issue date it was an ordinary software company with volatility of about 27%. The MSTR leg there reflects a high-beta technology stock in the COVID crisis rather than a crypto exposure, and the 2020 results should be read accordingly. A future version could separate the pre- and post-pivot periods or substitute an underlying that already carried crypto exposure in 2020.

11.2. Methodological Limitations

ERC on naked-asset volatility. The ERC weights are set from naked-asset volatilities, but the option overlays change each leg's risk profile — the covered call truncates the upside, the collar both sides — so the 'equal risk' condition is met only to first order, and the collared IBIT leg in particular contributes less true risk than the calculated figure. An iterative second-order decomposition, using effective volatilities and correlations computed from the Monte Carlo return distribution, is the natural refinement.

Look-ahead bias. Discussed in Section 8.3: the backtest remains a single realised path — a concrete illustration of behaviour, not a promise of performance — which is why the multi-scenario Monte Carlo analysis, and not the backtest, carries the main argument.

Parameter choice in the sensitivity tests. The sensitivity tests compare pre-chosen values of the DIP weight and floor level rather than solving a constrained optimisation, so the preferred configuration is the best among the candidates tested, not a demonstrated global optimum; a finer multi-parameter grid search would strengthen this claim.

Fixed parameters. Both pricing and simulation hold volatilities, correlations, and the risk-free rate constant over the three-year life. This understates tail risk in extreme co-movement scenarios, where volatilities and correlations rise together; a stochasticvolatility model or a time-varying-correlation copula would capture these effects more accurately.

European exercise in the tail hedge. The three DIP tranches are European and settle only at maturity, so the hedge pays on a downturn still depressed on the settlement date but nothing on a crash that recovers before it — the March 2020 pattern. It therefore protects against a persistent bear market of the 2000–2002 or 2008 type rather than a flash crash: a deliberate choice made for cost, but a real gap. American or Bermudan exercise, or a rolling series of shorter-dated puts, would close it at a materially higher premium.

11.3. Scope and Future Research

The issuer's side is not examined. The analysis is conducted entirely from the investor's side; it does not address embedded fees, dynamic-hedging costs, capital usage, or the issuer's profit-and-loss, and therefore cannot establish whether an issuer would offer the product on these terms. A full issuance-economics model is left to future work.

Beyond the improvements above, future research could re-estimate the analysis on IBIT's real trading history once it is long enough, model the issuer's side, extend the buy-and-hold structure to one that rebalances, and explore alternative underlyings, floor structures — including gradually released floors — and other forms of upside participation.

References

Baur, D. G. and Lucey, B. M. (2010). Is Gold a Hedge or a Safe Haven? An Analysis of Stocks, Bonds and Gold. *Financial Review* 45 (2), 217-229.

Baur, D. G. and McDermott, T. K. (2010). Is Gold a Safe Haven? International Evidence. *Journal of Banking & Finance* 34 (8), 1886-1898.

Bhansali, V. and Davis, J. (2010). Offensive Risk Management II: The Case for Active Tail Hedging. *The Journal of Portfolio Management* 37 (1), 78-91.

Black, F. and Scholes, M. (1973). The Pricing of Options and Corporate Liabilities. *Journal of Political Economy* 81 (3), 637-654.

Bouri, E., Molnár, P., Azzi, G., Roubaud, D. and Hagfors, L. I. (2017). On the Hedge and Safe Haven Properties of Bitcoin: Is It Really More Than a Diversifier? *Finance Research Letters* 20, 192-198.

Broadie, M., Glasserman, P. and Kou, S. (1997). A Continuity Correction for Discrete Barrier Options. *Mathematical Finance* 7 (4), 325-349.

Chemkha, R., BenSaïda, A., Ghorbel, A. and Tayachi, T. (2021). Hedge and Safe Haven Properties during COVID-19: Evidence from Bitcoin and Gold. *The Quarterly Review of Economics and Finance* 82, 71-85.

Choi, S. and Shin, J. (2022). Bitcoin: An Inflation Hedge but Not a Safe Haven. *Finance Research Letters* 46, 102379.

Cox, J. C., Ross, S. A. and Rubinstein, M. (1979). Option Pricing: A Simplified Approach. *Journal of Financial Economics* 7 (3), 229-263.

Glasserman, P. (2004). *Monte Carlo Methods in Financial Engineering*. New York: Springer.

Haug, E. G. (2007). *The Complete Guide to Option Pricing Formulas*, 2nd Ed. New York: McGraw-Hill.

Long, S., Pei, H., Tian, H. and Lang, K. (2021). Can Both Bitcoin and Gold Serve as Safe-Haven Assets? A Comparative Analysis Based on the NARDL Model. *International Review of Financial Analysis* 78, 101914.

Maillard, S., Roncalli, T. and Teïletche, J. (2010). The Properties of Equally Weighted Risk Contribution Portfolios. *The Journal of Portfolio Management* 36 (4), 60-70.

Reiner, E. and Rubinstein, M. (1991). Breaking Down the Barriers. *Risk* 4 (8), 28-35.

Stoimenov, P. A. and Wilkens, S. (2005). Are Structured Products 'Fairly' Priced? An Analysis of the German Market for Equity-Linked Instruments. *Journal of Banking & Finance* 29 (12), 2971-2993.

Appendix A — Data and Sources

All market data for the four underlying assets — the SPDR Gold Shares ETF (GLD), the iShares Bitcoin Trust ETF (IBIT), MicroStrategy (MSTR), and the S&P 500 index (SPX) — were obtained from Bloomberg. This covers the issue-date spot prices, dividend yields, and the daily series used to estimate the trailing realised volatilities, together with the three-year US Treasury yield used as the risk-free rate.

The IBIT spot ETF did not begin trading until January 2024, so no real IBIT price history exists for the issue dates used in this work (January 2023, and the earlier robustness dates in 2021 and 2020). For the period before listing, a synthetic IBIT series was constructed from the spot price of Bitcoin, which was obtained from Yahoo Finance. After confirming that Bitcoin's daily log returns track those of IBIT over the window in which both trade, the IBIT series was extended backwards day by day from its price on the first issue date, applying the Bitcoin daily returns. The 65.71% volatility used for the IBIT leg is therefore estimated from this synthetic series; the limitations of this proxy are discussed in Section 11.1.

Appendix B — Core Python Code

The Python code below was produced with the AI assistance described in the Declaration on the Use of Generative AI at the front of this work. The author specified the model structure and parameters, reviewed the implementation, and validated all numerical outputs against the pricing workbook. The listings are the core pricing and simulation routines; input/output handling and helper utilities are omitted for brevity.

Listing B.1 — Monte Carlo engine core: correlated GBM paths, barrier monitoring, product mark-to-market.

```
import numpy as np

# SIGMA, L (Cholesky factor of the correlation matrix), S0, DT, N_STEPS,
# R, Q and the strike/barrier constants are the issue-date calibration;
# N[...] are the per-leg unit counts and fp is the floor payoff (70).
rng = np.random.default_rng(seed)
log_drift = (mu - Q - 0.5 * SIGMA**2) * DT
log_vol   = SIGMA * np.sqrt(DT)

for start in range(0, n_paths, batch_size):
    nb   = min(batch_size, n_paths - start)
    logS = np.tile(np.log(S0), (nb, 1))
    spx_min = np.full(nb, np.inf)
    for step in range(1, N_STEPS + 1):
        Z    = rng.standard_normal((nb, 4)) @ L.T   # impose correlation
        logS = logS + log_drift + log_vol * Z       # GBM increment
        S    = np.exp(logS)
        spx_min = np.minimum(spx_min, S[:, 3])       # track S&P minimum
        tau  = max(T - step * DT, 1e-6)
        gld  = S[:, 0] - bs_call(S[:, 0], K_CALL_GLD, sig[0], tau, R, 0)
        ibit = (S[:, 1] + bs_put(S[:, 1], K_PUT_IBIT, sig[1], tau, R, 0)
                        - bs_call(S[:, 1], K_CALL_IBIT, sig[1], tau, R, 0))
        mstr = (bs_call(S[:, 2], K_LONG_MSTR,  sig[2], tau, R, 0)
                - bs_call(S[:, 2], K_SHORT_MSTR, sig[2], tau, R, 0))
        dip  = np.zeros(nb)
        for B, K in DIP_SPECS:             # layered down-and-in puts
            dip += down_in_put_vec(S[:, 3], K, B, sig[3], tau,
                                   R, Q_SPX, spx_min <= B)
        mv = (fp * np.exp(-R * tau)
              + N['GLD'] * gld + N['IBIT'] * ibit
              + N['MSTR'] * mstr + N['DIP'] * dip)     # product value
```

Appendix C — Additional Figures and Tables

This appendix collects the auxiliary terms of the barrier-option formula used in Chapter 5, together with the supporting figures and full sensitivity tables referenced in Chapters 7–9.

C.1. Auxiliary Terms of the Reiner–Rubinstein Down-and-In Put

The closed-form value in equation (5.4) uses three auxiliary terms. Writing S for the spot, K for the strike, B for the barrier, σ for the volatility, T for the time to maturity, and $\lambda = (r - q + \sigma^2/2)/\sigma^2$ as defined in equation (5.3), they are:

$$x_1 = \ln(S/K) / (\sigma\sqrt{T}) + \lambda\sigma\sqrt{T}$$
$$y = \ln(B^2/(SK)) / (\sigma\sqrt{T}) + \lambda\sigma\sqrt{T}$$
$$y_1 = \ln(B/S) / (\sigma\sqrt{T}) + \lambda\sigma\sqrt{T}$$

The term x_1 is the familiar Black–Scholes d_1 with λ in place of the usual drift adjustment, so the first two terms of equation (5.4) reproduce an ordinary European put. The terms y and y_1 are its reflections about the barrier, and they are what generate the $(B/S)^{(2\lambda)}$ factors: these carry the probability mass of the paths that reach B and therefore knock the option in. The expressions follow Haug (2007) for the case in which the barrier lies at or above the strike, which is the configuration of all three tranches used here; when $B < K$ the formula takes a different form, which is not needed in this work.

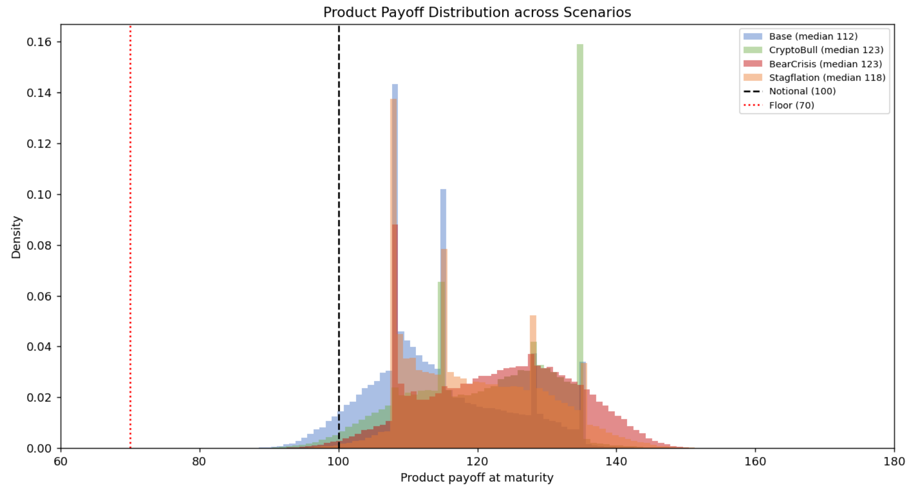

FIGURE 3 – Product payoff distribution across the five scenarios.

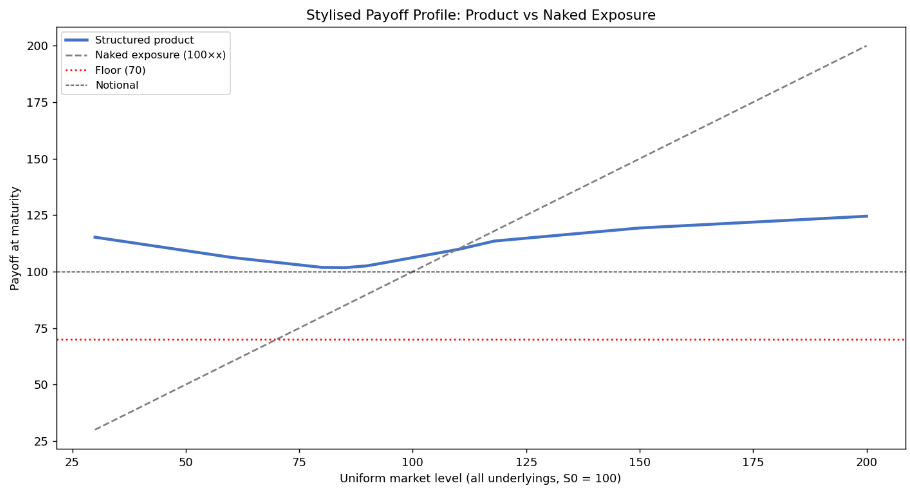

FIGURE 4 – Stylised payoff profile: product versus naked exposure under a uniform market move.

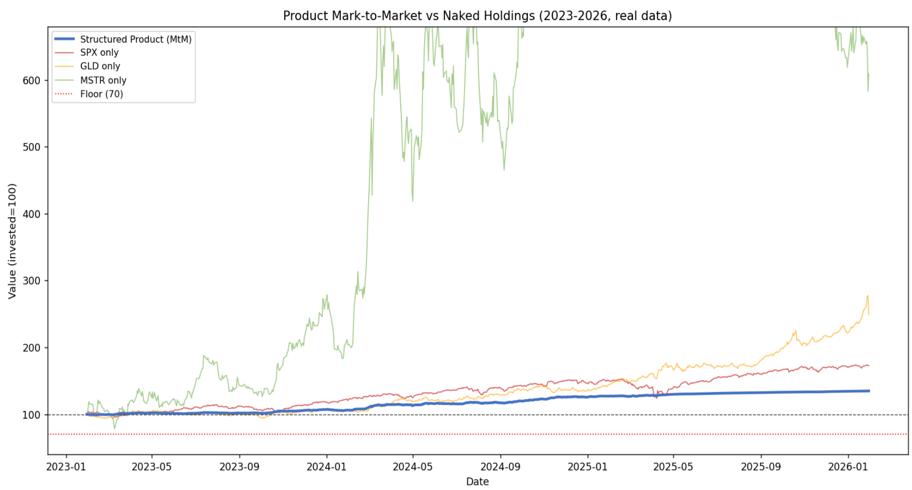

FIGURE 5 – Product mark-to-market versus naked holdings (2023–2026, real data).

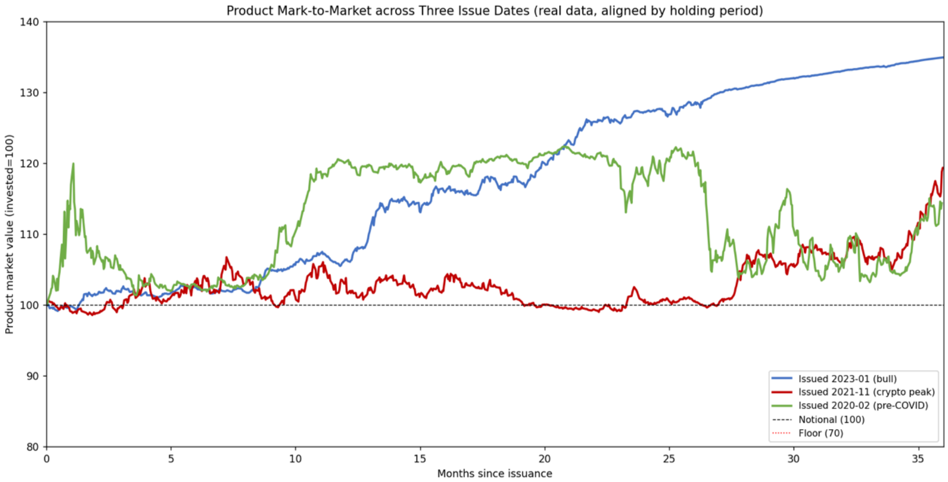

FIGURE 6 – Product mark-to-market across three issue dates (aligned by holding period).

TABLE XIII
SENSITIVITY TO THE DIP WEIGHT

DIP weight	Base E[ann]	Base Sharpe	Base P(loss)	Base vol	CryptoBull E[ann]	BearCrisis E[ann]	BearCrisis worst
10%	4.35%	+0.084	5.1%	4.80%	6.71%	6.98%	87.2
20%	3.69%	−0.042	11.5%	6.12%	5.84%	10.15%	85.2
30%	2.97%	−0.123	33.7%	7.97%	4.92%	13.03%	81.5

Source: Author's calculations, 100,000 simulated paths per case. E[ann] and vol are annualised; BearCrisis worst is the lowest terminal value across paths, per 100 units of notional.

TABLE XIV
SENSITIVITY TO THE FLOOR LEVEL

Floor	Maturity floor	Floor cost (PV)	Risky sleeve	GLD	IBIT	DIP	MSTR
50%	50.00	44.42	55.58	38.34	8.90	5.56	2.78
70%	70.00	62.19	37.81	26.08	6.06	3.78	1.89
90%	90.00	79.95	20.05	13.83	3.21	2.00	1.00

Source: Author's calculations. The 70% row is the floor used in the product. The risky sleeve and the four leg budgets are stated per 100 units of notional.

TABLE XV
SENSITIVITY TO THE IBIT VOLATILITY

σ (IBIT)	Long put	Short call	Collar unit value	N (units)	Δ unit value	Δ N
65.71% (synthetic)	1.8538	4.6694	10.5192	0.575750	—	—
50.00% (real)	1.0917	3.2567	11.1698	0.542216	+6.18%	−5.82%

Source: Author's calculations. The synthetic case (65.71%) is the one used in the rest of the work; the last two columns are measured against it. $S_0 = 13.33$, put strike at 60%, call strike at 150%, IBIT budget fixed at 6.06 per 100 notional.

TABLE XVI
UNDERLYING PRICES AT ISSUANCE AND MATURITY

Asset	S0 (30/01/2023)	S_T (30/01/2026)	Change	Note
GLD	178.76	444.95	148.9%	
IBIT	13.33	47.49	256.1%	synthetic
MSTR	24.57	149.71	509.3%	
SPX	4,017.77	6,939.03	72.7%	

Source: Author's calculations from historical data. The IBIT issue-date price is taken from the synthetic series.

TABLE XVII
PRODUCT VERSUS NAKED HOLDINGS (SAME PERIOD)

Strategy	Terminal value	Annualised return
Structured product	134.91	10.50%
All in GLD	248.91	35.52%
All in SPX	172.71	19.98%
All in MSTR	609.35	82.65%
Four assets equal (25%)	346.77	51.36%

Source: Author's calculations over 30/01/2023 to 30/01/2026, on an invested base of 100.