**QUANT RESEARCH**

# Option Pricing in Metal Markets: A Practical Framework Combining Theory and Trade Strategy

This report explores the mathematical frameworks used to price metal commodity options and applies them to a real-world trade in aluminium. As Nassim Taleb wrote, *“Options, by their very nature, are convex; they benefit from volatility, from disorder, from uncertainty.”*

**Martin Carrasco**
martintuan.carrasco@alumni.esade.edu
*Team Lead & Research Analyst, ESFS*
**Francesco Palma**
francesco.palma@alumni.esade.edu
*Research Analyst, ESFS*
**Angel Brendel**
angelpaulnicolasmarcemile.brendel@alumni.esade.edu
*Research Analyst, ESFS*

# CONTENTS PAGE


| | |
|---|---|
| **Executive Summary** | **3** |
| **Theory for Option Pricing** | **4 - 8** |
| - Fundamentals of Option Pricing | 4 - 5 |
| - Black-Scholes Theory | 5 - 6 |
| - Black 76 Theory | 7 |
| - Black-Scholes Errors | 7 |
| - Implied Volatility Surface | 8 |
| **Practical Applications** | **9 - 13** |
| - Variables Used Summary | 9 |
| - Black 76 Application | 9 - 10 |
| - Schwartz One-Factor Application | 10 - 11 |
| - Heston-Model Application | 12 - 13 |
| **Trade Idea on Aluminium** | **14 - 18** |
| - Trade Idea Summary | 14 |
| - Macroeconomic Analysis | 15 |
| - Implied Volatility Analysis | 16 |
| - Skew-Adjusted Implied Volatility Analysis | 17 |
| - Monte Carlo Simulation Analysis | 18 |
| **Conclusion** | **19** |
| **Appendix & Bibliography** | **20** |

# Executive Summary

In the past two decades, the use of options in metal commodity markets has grown significantly. Traders, hedgers and investors seek to manage risk, speculate on price movements and build strategies. These options, often written on metal futures contracts, offer a broad range of possibilities. However, accurately pricing them remains a challenge. Indeed, unlike equities or fixed-income assets, metal prices have specific features that require some adjustments to classical financial models to ensure pricing accuracy and practical relevance.

This paper aims to serve two main purposes.

First, it provides a concise and structured review of the main mathematical frameworks used for pricing options on metal commodities. Starting from the foundational Black 76 model, we progress to more commodity-specific models such as the Schwartz family of mean-reverting processes and stochastic volatility models like Heston. Each model is discussed in terms of assumptions, practicality and limitations when applied to metal markets.

Second, we bridge theory with practical application by proposing a live trade idea in a selected metal market. Using current market data, we analyse the term structure of implied volatility and design a structured option trade. The trade is defined in terms of entry conditions, risk management rules, expected payoff scenarios and potential limitations.

We will first review the fundamentals of option pricing. Then, we will cover key theoretical models and their suitability for metals. Afterwards, we will present a structured trade idea based on market conditions. Finally, we will interpret the results, conclude the findings and answer our research question: How well do traditional models price metal options and how can they inform trading decisions?

# Theory for Option Pricing

An option is a contract that constitutes the right to buy (call option) or to sell (put option) an asset (the underlying asset in the contract) at a fixed price (strike price) and at/by a set date.

To price options, we consider the Current Price of the Underlying ($S_0$), Strike Price ($K$), Time to Maturity ($T$), Volatility of the Underlying (s), Risk-free Rate ($r$) and possible Dividends ($q$).

There are several position types:

| Position | Right / Obligation | Action on the Underlying |
|---|---|---|
| ***Long Call*** | Right to buy the asset | Buy the asset |
| ***Short Call*** | Obligation to sell the asset if exercised | Sell the asset |
| ***Long Put*** | Right to sell the asset | Sell the asset |
| ***Short Put*** | Obligation to buy the asset if exercised | Buy the asset |

Figure 1: Call Option Payoffs Comparison

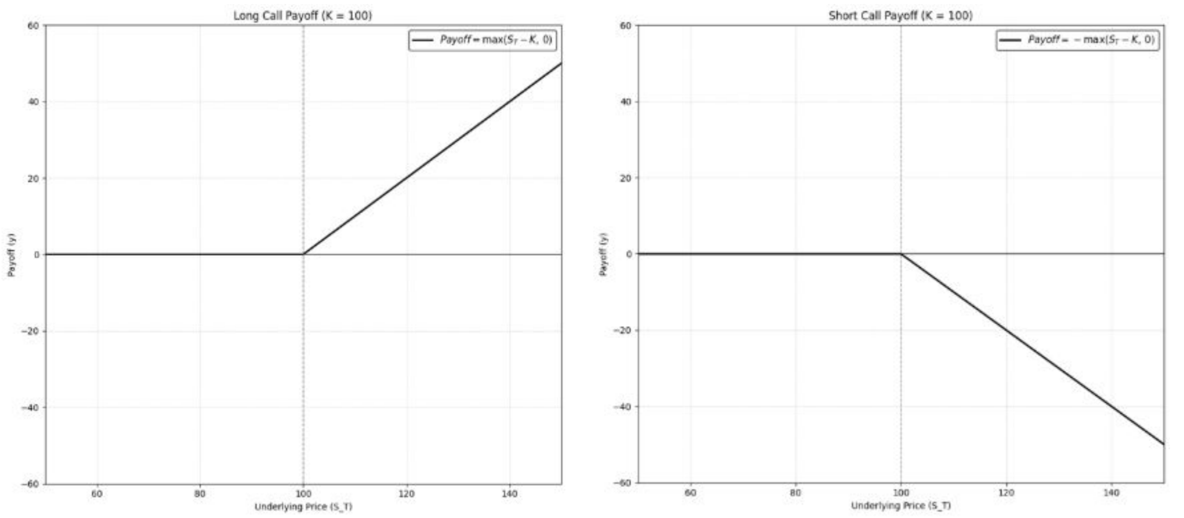

# Theory for Option Pricing

Figure 2: Put Option Payoffs Comparison

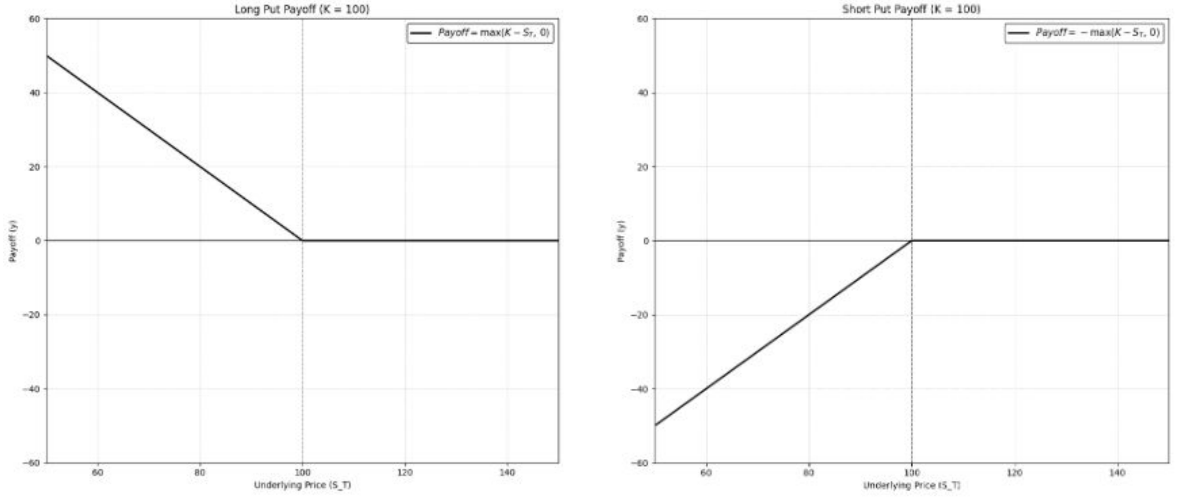


The pricing of options, whether European or American, lies at the intersection of financial economics, stochastic calculus and arbitrage theory, as both types of contracts derive their value from the underlying asset's uncertain future price. However, they differ fundamentally in the right to exercise. Indeed, a European option can be exercised only at maturity, while an American option can be exercised at any point before expiration. This is a difference that introduces conceptual and mathematical distinctions in their valuation frameworks.

**The Black-Scholes Theory for Option Pricing**

The Black-Scholes-Merton model, developed in 1973, provides a mathematical framework to determine the fair price of a European option under a set of simplifying assumptions.

The model assumes that the underlying asset follows a Geometric Brownian Motion (GBM) with constant volatility ($\sigma$), constant risk-free rate ($r$) and the Wiener process $W_t$ that represents the random noise.

$$dS = \mu S_t dt + \sigma S_t dW_t$$

Where $\mu$ is the drift rate of the process (the average growth rate of $S_t$).

# Theory for Option Pricing

Under these assumptions, a continuously adjusted portfolio consisting of the underlying asset and a risk-free bond can perfectly replicate the option's payoff, meaning that the two must have the same price (no-arbitrage principle).

The GBM reflects the idea that stock prices evolve continuously and randomly over time, with both a predictable trend and an unpredictable component (the noise).

The Black-Scholes formula then gives the theoretical price of a European call as:

$$C = S_0 e^{-qT} N(d_1) - K e^{-rT} N(d_2)$$

and for a put option:

$$P = K e^{-rT} N(-d_2) - S_0 e^{-qT} N(-d_1)$$

where:

$$d_1 = \frac{ln(\frac{S_0}{K}) + (r - q + \frac{\sigma^2}{2})}{\sigma\sqrt{T}} \quad , \quad d_2 = d_1 - \sigma\sqrt{T}$$

and $N(d)$ is the cumulative normal distribution.

Each term has a clear interpretation:

- $S_0 e^{-qT} N(d_1)$ represents the present value of holding the asset, adjusted for dividends and the probability the option ends up in the money.
- $K e^{-rT} N(d_2)$ is the present value of paying the strike price, weighted by the same probability under a risk-neutral world.

An option is said to be in the money (ITM) when exercising it immediately would generate a positive payoff, which occurs when the underlying price exceeds the strike for a call option or when it falls below the strike for a put.
It is at the money (ATM) when the underlying price is approximately equal to the strike.
It is out of the money (OTM) when exercising it would have no value, that is when the underlying price is below the strike for a call or above it for a put.

These distinctions matter because they determine how likely the option is to end profitably, which is exactly what the terms $N(d_1)$ and $N(d_2)$ in the Black-Scholes model represent: probabilities that the option will finish in the money and measures of its sensitivity to changes in the underlying price.

# Theory for Option Pricing

**Black-76**

Commodities involve storage costs (the expenses associated with holding a physical commodity over time) and convenience yields (the non-monetary benefits of possessing the commodity, ex: a refinery holding crude oil to ensure production continuity). Therefore, traders typically use futures contracts rather than spot prices to value options. Futures prices already reflect expectations about interest rates, storage and the benefits of holding the physical good. Hence, these effects are managed implicitly through the futures curve.

In a futures contract, the underlying already reflects the cost of carry, so there is no discount of future spot prices.

$$C = e^{-rT}[F_0 N(d_1) - KN(d_2)]$$

However, commodity markets exhibit specific features that challenge the standard model's assumptions. Unlike financial assets, commodities are heavily influenced by seasonal cycles (ex: heating demand in winter or crop harvest periods in summer, which lead to pronounced fluctuations in both prices and volatilities).

These effects are often amplified by supply shocks, such as weather disruptions, geopolitical tensions or inventory shortages, which can alter expected future availability. As a result, the implied volatility surface for commodity options tends to display distinct and often steeper patterns across maturities compared to equities or interest rate products, with evident seasonal skews and term structures.

Consequently, while the theoretical framework of models like Black-76 remains valid, interpreting their outputs in the context of commodity markets requires recognizing these unique dynamics, particularly how storage constraints, demand cycles and supply uncertainty shape the behaviour of prices and implied volatilities.

**Black-Scholes Errors**

While the Black-Scholes-Merton model assumes a constant volatility across all strikes and maturities, in real markets, this assumption does not hold.

When option prices observed in the market are inverted through the Black-Scholes formula to obtain the implied volatility (volatility consistent with each option's market price), the resulting values differ from those predicted by the model.

When implied volatilities are plotted across strikes and maturities, they produce what is known as the implied volatility surface. Ideally, if the Black-Scholes assumptions were perfectly valid, this surface would be flat, indicating a single volatility level for all options.

# Theory for Option Pricing

However, empirical evidence shows that the surface exhibits distinct patterns, such as the volatility smile or skew. Implied volatility tends to be higher for options that are deep in-the-money or out-of-the-money, reflecting asymmetric expectations about large price movements. Another typical characteristic of the pattern is the term structure, since implied volatility also varies with time to maturity, often decreasing for shorter-dated options and increasing for longer horizons.

The following figure presents a simulated volatility surface generated by inverting Black-Scholes prices imputed under a non-constant volatility function.

Figure 3: Simulated Implied Volatility Surface (under a non-constant volatility function)

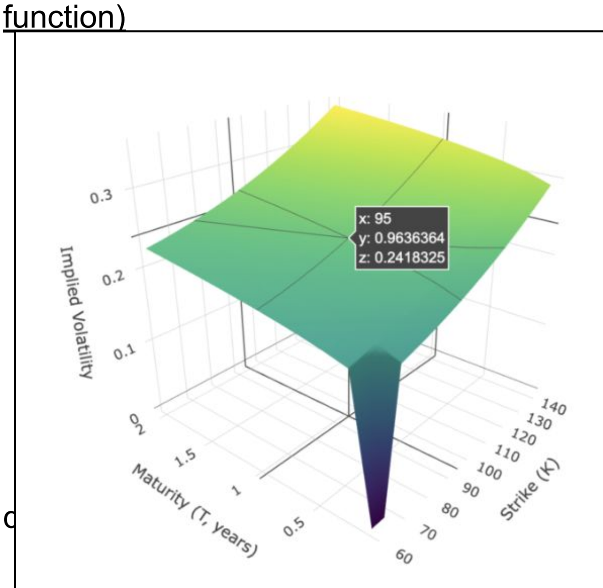


controlled environment (with lated randomly but under the ptions), the implied volatilities ss strikes and maturities, g how sensitive the model is to from its constant-σ on. The curvature illustrates the smile" across strikes and the ture across maturities. This ates visually one of the model's itations: volatility is not nd state-

rent state or level of the g variable).

To capture the shapes observed in the volatility surface, modern approaches extend the original model by allowing volatility to vary either with the underlying price or over time.

Among the most prominent extensions are:

- Local volatility models (Dupire, Derman-Kani), where volatility is a deterministic function of both price and time.
- Stochastic volatility models (Heston, Hull-White), where volatility follows its own random process.
- Jump-diffusion models (Merton), which incorporate sudden price jumps to explain extreme option pricing behaviour.

These extensions preserve the arbitrage-free foundation of Black-Scholes but relax the constant-volatility assumption. Thus, it improves fit to data and allows a more accurate valuation of complex derivatives.

# Practical Applications

| Variable | Symbol | Value | Description |
|---|---|---|---|
| Initial futures price | | 2,855.5 | Starting price of the underlying futures contract |
| Strike price | | 2,950 | Exercise price of the option |
| Risk-free rate | | 0.038 | US 3-Month Treasury Bills |
| Mean reversion speed | | 4 | Speed at which variance reverts to its long-term mean *(Appendix)* |
| Long-run variance | | 0.05 | Average level around which variance fluctuates *(Appendix)* |
| Volatility of volatility | | 0.35 | Volatility of the variance process *(Appendix)* |
| Correlation | | -0.2 | Correlation between futures' price and variance innovations *(Appendix)* |
| Initial variance | | 0.0310 | |
| Time to maturity | | 0.1123 | Time in years until expiry (≈ 41 days) |
| Number of paths | | 1,000 | Monte Carlo simulation paths *(Appendix)* |
| Time steps | | 500 | Discretization of each path *(Appendix)* |

Contract: *ALEH6 Call (Source: Bloomberg)*

**Black 76**

The first model implemented in Python corresponds precisely to this framework (in reference to theory as it follows it), adapted to the case where the underlying is a futures contract rather than a spot asset.

The pricing follows the same Black-Scholes logic, but adapted as "Black-76 formulation", where the spot price ($S_0$) is replaced by the current futures price ($F_0$). Since futures prices already incorporate the cost of carry (interest rates, storage costs and convenience yields), there is no need to discount future spot prices in the formula.

# Practical Applications

As a result, under the risk-neutral measure, the expected drift of the futures price is 0, which simplifies the expression for $(d_1)$ and removes the term $(r - q)$ that appears when pricing options on spot assets. In this setup, the only discounting that remains is the factor $(e^{-rT})$, which adjusts for the time value of money between the valuation date and maturity.

The numerical implementation uses the same parameters as those assumed in the theoretical model. The numerical evaluation of the model gives a theoretical call option price of 34.5559, representing the fair value implied by the Black-76 formulation.

Figure 4: Option Value, Surface on $ALEH6\ Call$

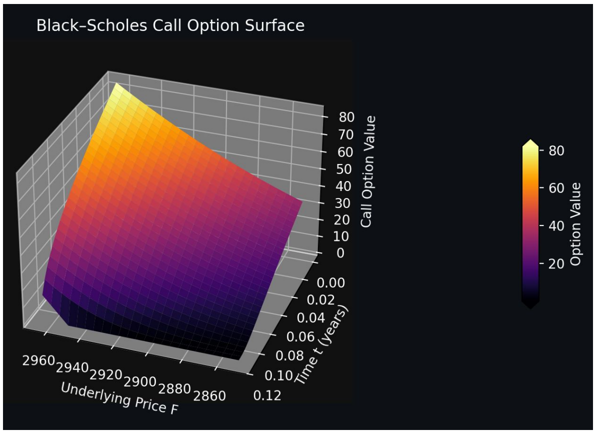


The figure above illustrates the surface of call option values as a function of the underlying price $F$ and the time to maturity $t$. The surface shows how the option value increases with higher underlying prices and decreases as time passes, reflecting both the sensitivity of the option to price movements $(\delta)$ and its progressive time decay $(\theta)$.

**Schwartz Models**

The Schwartz one-factor model (1997) was originally developed to describe how commodity prices (ex: oil, gas or metals) evolve over time. The model starts with the idea that a commodity's spot price (its current market price) changes randomly. However, it also depends on an economic factor called the convenience yield, which represents the practical benefit of holding the physical commodity (ex: ensuring supply or avoiding storage costs).

From this relationship, the futures price (the price agreed upon today for delivery in the future) can be expressed using the cost-of-carry relation, which links the spot and futures markets.

# Practical Applications

While the Schwartz one-factor model assumes that the spot price follows a mean-reverting process under the real (physical) measure, this changes once we move to the risk-neutral framework used for pricing. Under the risk-neutral measure, the corresponding futures price can be shown to follow a lognormal distribution (positive, skewed, exponential-like distribution) similar to the assumption in the Black-76 model. This relationship shows why Black's formula can still be used to price options on futures, even though the underlying spot price dynamics in the Schwartz model are more complex.

However, when the option can be exercised early (as in American options), there is no exact mathematical solution. In this case, numerical methods like the Least Squares Monte Carlo (LSM) technique must be used to estimate the price. Importantly, the LSM method is not part of the original Schwartz model, but rather a numerical tool that can be applied to any model when early exercise is possible. Finally, extending this approach to the two- or three-factor versions of the Schwartz model is not relevant here, since those formulations include explicit drift terms for the spot price. In the case of futures, the drift is already embedded in the futures price itself under the risk-neutral measure, making such extensions unnecessary for our purpose.

The parameters used in this model are consistent with those of the previous one.

The numerical results show that the estimated value of the American call option on futures under these parameters is approximately 34.9796.

<u>Figure 5: Montecarlo simulation on *ALEH6 Call under Schwartz One-Factor Model*</u>

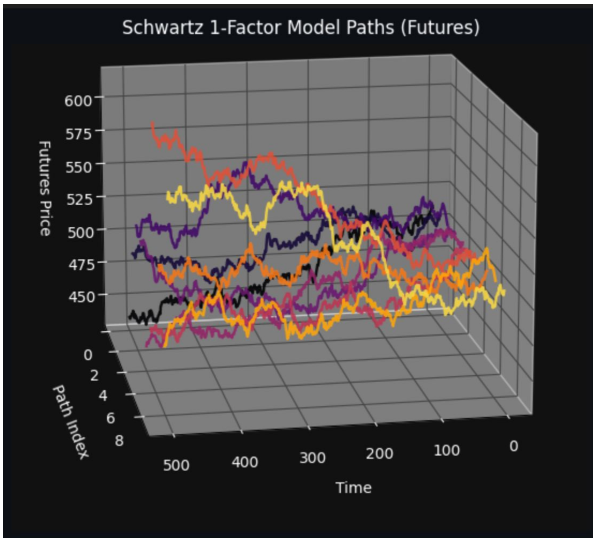


The simulated paths show random, driftless evolution of futures prices; option value reflects the flexibility of early exercise captured by the LSM Monte Carlo method. The figure displays a selection of simulated futures price paths generated under the Schwartz one-factor model. Each trajectory represents a possible evolution of the futures price through time, reflecting random market movements governed by volatility. As expected, the prices fluctuate around their initial value, illustrating the absence of drift under the risk-neutral measure.

# Practical Applications

**Heston-Model**

The final model implemented extends the previous frameworks by introducing stochastic volatility through the Heston model, adapted here for futures contracts. While the Black-76 and Schwartz one-factor models assume a constant volatility over time, the Heston model allows volatility itself to evolve randomly according to a mean-reverting process. This captures one of the most important empirical features of real markets: volatility is not constant but tends to fluctuate around a long-term average and exhibits correlation with price movements.

In this model, both the futures price and its variance are simulated jointly, with correlated random shocks to reflect the “leverage effect” (the tendency for volatility to rise when prices fall). Under the futures measure, the drift of the futures price remains null, consistent with the property assumed in the previous models.

The simulation produces an estimated American call option value of approximately 34.3312 under the Heston dynamics.

Figure 6: Montecarlo simulation on *ALEH6 Call under Heston Model*

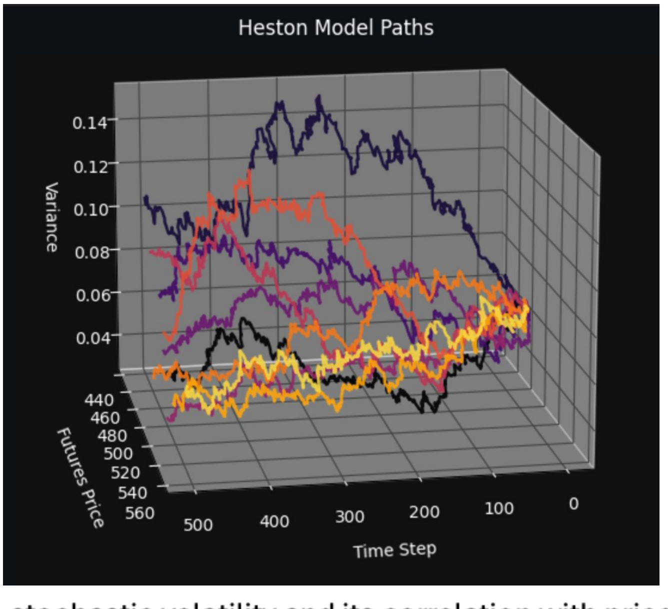


The 3D plot displays a subset of simulated trajectories showing the joint evolution of the futures price and its instantaneous variance. The coloured paths highlight how volatility changes over time, producing heteroscedastic (uneven) price movements that differ significantly from the smoother diffusion of the Schwartz one-factor model. This richer dynamic allows the Heston framework to reproduce features observed in real commodity markets, such as volatility clustering and the asymmetric relationship between price and variance. This model captures stochastic volatility and its correlation with price movements, showing how variable volatility leads to more realistic price dynamics and option values slightly higher than those under constant volatility assumptions.

# Practical Applications

**Comparison and Critical Assessment**

The Black-76, Schwartz one-factor and Heston models trace a clear evolution in option pricing theory.

The Black-76 model is the most straightforward and widely used foundation. It assumes prices follow a random process with constant volatility and no drift under the risk-neutral measure, leading to a clean, closed-form solution for European options. Its strength lies in clarity and ease of use, but its assumption of fixed volatility limits its ability to capture real market dynamics such as changing volatility or term structure effects.

The Schwartz one-factor model adds an economic layer tailored to commodities. It connects spot and futures prices through the cost of carry and introduces the convenience yield, which reflects the benefit of physically holding the commodity. This makes it more realistic for metal markets. Yet, similarly to Black-76, it still assumes constant volatility and therefore misses the variability often seen in practice.

The Heston model advances further by allowing volatility itself to be stochastic, fluctuating over time and potentially correlated with price changes. This feature enables it to replicate real market patterns such as the volatility smile and the leverage effect, offering a more faithful representation of observed behaviour. However, this realism comes at the cost of higher complexity and more demanding calibration.

In short, the three models represent a continuum: Black-76 provides analytical clarity, Schwartz adds economic meaning and Heston brings statistical depth. The choice among them depends on the balance one seeks between simplicity, interpretability and realism in capturing how markets truly behave.

# Trade Idea on Aluminium

| Trade Idea Summary | |
|---|---|
| **Contract** | **ALEK6** |
| **Expiry Month** | February 2026 |
| **Days to Expiry** | 117 |
| **Futures Price ($F_0$)** | 2,877.75 USD/tonne |
| **Strike (K)** | 2,950 USD/tonne |
| **Option Type** | Call |
| **Contract Size** | 25 metric tonnes |
| **Risk-Free Rate** | 4.02% |
| **Market Price** | 106.50 USD |
| **Implied Volatility** | 25.01% |
| **Delta** | 0.453 |
| **Gamma** | 0.000961 |
| **Vega** | 638.19 |
| **Theoretical Price** | 129.21 USD |
| **Model Mispricing (dTheo)** | –22.71 USD |
| **Break-even Level** | 3,056.50 USD/tonne |

The aluminium market is approaching a pivotal structural shift that supports a long call position on the February 2026 LME Aluminium contract.

The ESFS Quant Team selected a strike near USD 2,950/tonne, which would provide inexpensive exposure to a tightening physical market expected to transition from surplus in 2025 to deficit in 2026, yet undertaking a medium-high level of risk. The break-even price of USD 3,056.5/t is obtained by adding the strike (2,950) to the option premium (106.50), implying the underlying must rise ~6.2 % from current levels (2,877.75 USD/t) for the position to turn profitable.

The payoff diagram shows the profit and loss profile of the February 2026 aluminium call option with a strike of $2,950/t, with the position incurring a maximum loss equal to the premium paid ($106.50/t) if prices remain below the strike. On the other hand, profits increase linearly beyond the break-even level of $3,056.50/t. The flat region below the strike represents the option expiring worthless, whereas the upward-sloping section highlights the upside potential as aluminium futures rise.

# Trade Idea on Aluminium

Figure 7: Payoff at Expiry


In terms of Macro context, China's 45 Mt output cap and slow capacity restarts outside the country keep supply tight. European smelters remain idle after the energy crisis, US production has fallen to just 1% of global output and new projects in Indonesia and India face delays.

Additionally, inventories on both London Metal Exchange and Shanghai Futures Exchange are at multi-year lows (LME down ~50 % YoY), while US tariffs and a +177 % Midwest premium further restrict physical availability. Analysts from Bank of America and UniCredit see prices approaching USD 3,000/t as a result.

From a market structure standpoint, implied volatility remains exceptionally low relative to historical norms, offering "cheap" optionality with limited time decay given the tenor and stable rate environment. With spot prices near USD 2,750-2,800, a move to USD 3,000 would imply a 7-9% upside in the underlying, providing a favourable risk/reward scenario. Moreover, the model mispricing (dTheo = -22.7 USD) indicates the call trades below its theoretical fair value, suggesting further attractive entry levels for the wanted exposure.

# Trade Idea on Aluminium

Beyond the absolute level of implied volatility, it is worth examining how volatility changes across strikes, a pattern known as the volatility skew or smile. This skew reflects where investors are placing their hedging or speculative demand: higher volatility at lower strikes usually signals demand for downside protection, while higher volatility at higher strikes reflects expectations of upside risk.

In this case, the implied volatility smiles across expiries reveal a clear upward-sloping term structure, signalling that the market is increasingly pricing in upside uncertainty for 2026. Shorter maturities display a flatter U-shaped pattern, while longer expiries such as the February 2026 contract show both higher volatility and steeper slopes across strikes. This suggests that traders are attributing greater probability to large price swings (and particularly upward moves) as the horizon extends. The chosen Feb 2026, 2950-strike call, positioned slightly out of the money, lies in the moderate region of the volatility curve where convexity is attractively priced, meaning that the option's price offers good value for its potential upside, offering strong payoff potential (convexity) for a relatively cheap cost.

Summarising, the volatility smiles indicate growing long-term bullish uncertainty, with upside risks increasingly priced into 2026 maturities.

Figure 8: Implied Volatility Surface on *ALEG6, ALEH6, ALEJ6, ALEK6 Call*

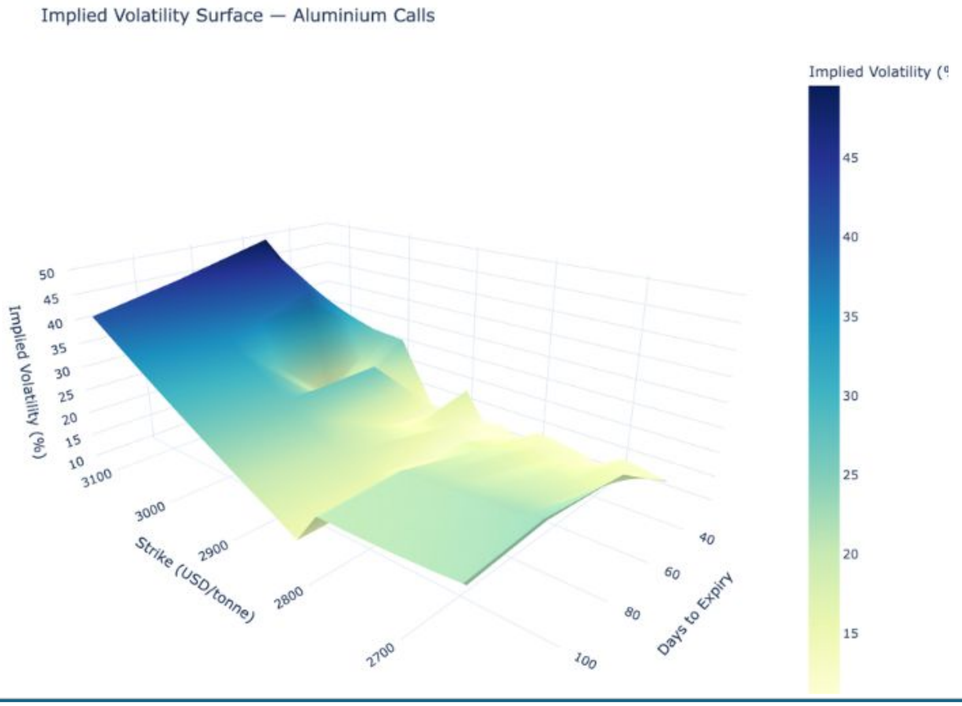

# Trade Idea on Aluminium

However, a single option can still be mispriced relative to this overall smile. To identify that, we fit a smooth curve of implied volatility against strike for each expiry and measure how far each option lies from the curve, obtaining its skew-adjusted residual. A positive residual means the option is expensive (the market is overpaying for volatility), while a negative residual means the option is cheap (volatility is underpriced).

In the chart below, the February 2026 call at $2,950 shows a residual of -2.88, meaning its implied volatility is roughly 2.9 percentage points lower than what the curve would predict for similar options. In plain terms, the market appears to be undervaluing volatility for this strike. This makes the option an efficient way to gain exposure to potential upside: it offers more convexity per dollar of premium, as profits would accelerate faster if aluminium prices rose beyond the break-even level.

Figure 9: Skew-Adjusted Implied Volatility (Residual vs Log-Moneyness Fit)

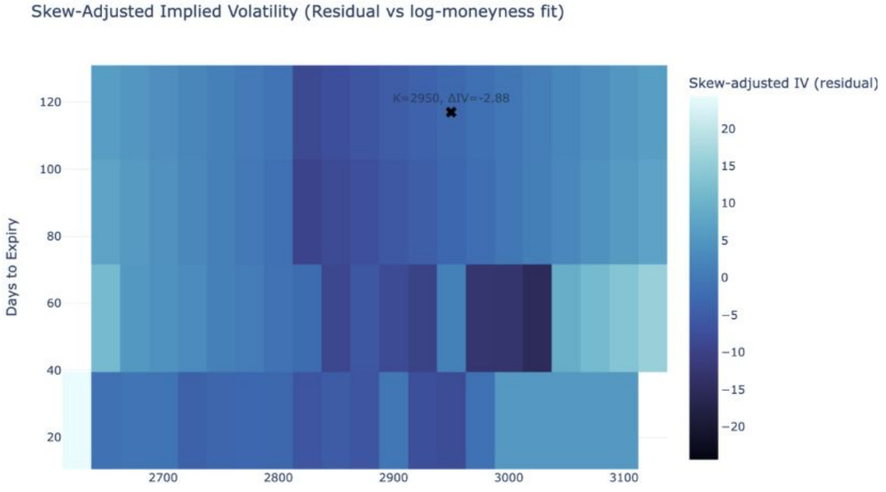


To complement the volatility surface study, a Monte Carlo simulation was performed to illustrate the potential range of aluminium futures outcomes until February 2026 under a geometric Brownian motion framework. Starting from USD 2,877.75, the model assumes an 8% annual drift and 20% annualised volatility, generating 300 potential price paths.

# Trade Idea on Aluminium

The simulated distribution converges around a mean terminal price of USD 2,950, almost exactly matching the strike, with roughly 50% of paths finishing in-the-money. This balanced probability underscores that the market-implied forward level is close to the strike, making the option's value highly sensitive to volatility and drift assumptions.

The analysis highlights the position's asymmetric payoff: the downside is capped at the premium paid, while the upside increases linearly beyond the break-even level of USD 3,056.5/t. This convexity provides a cost-efficient exposure to potential medium-term bullish developments in aluminium prices.

Figure 10: Monte Carlo Simulation on $ALEK6$ $(Feb$ $2026)$ $Calls$, with $K$ at 2,950

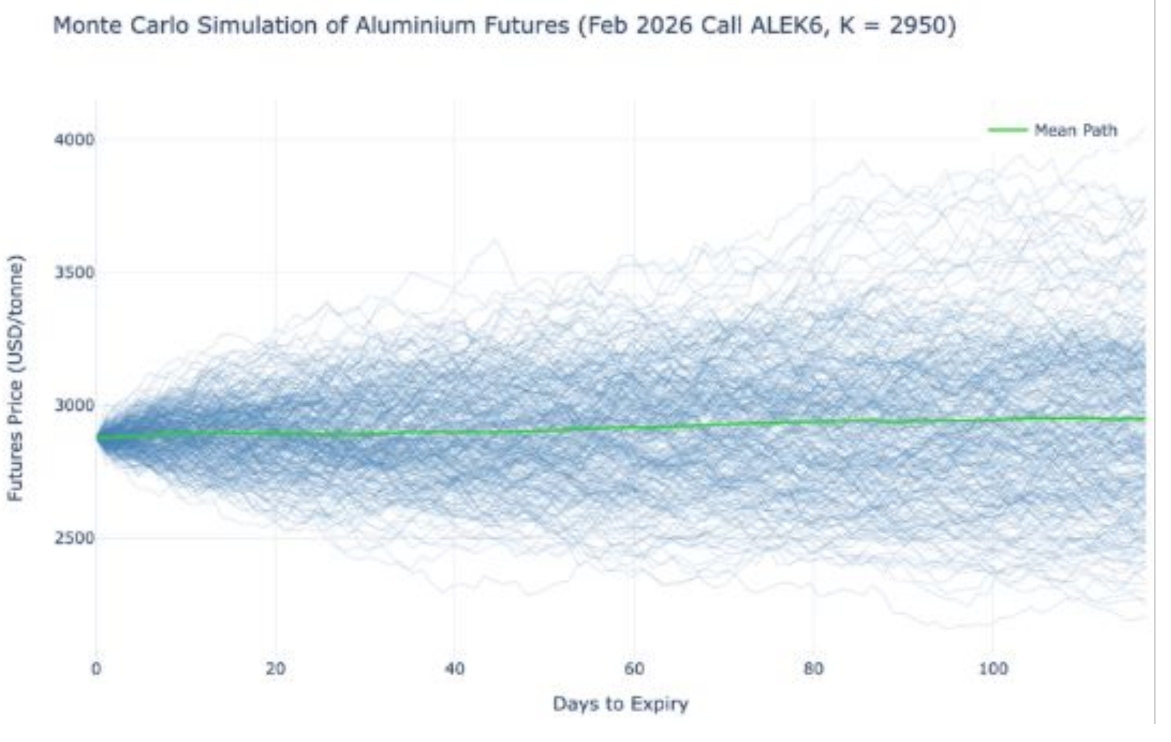


Overall, the combination of strong macro fundamentals, supply constraints and historically cheap implied volatility creates a compelling case for a long aluminium call position expiring in February 2026. In essence, this trade captures convex upside exposure to aluminium's structural revaluation at a time when the option market still discounts volatility, a classic asymmetric setup consistent with professional commodities option strategy frameworks.

# Conclusion

This paper has explored the theoretical and practical dimensions of option pricing in metal commodity markets, with an additional trade idea on aluminium. We began by reviewing the classical Black-Scholes framework and its commodity-specific adaptation with the Black-76 model. While these models offer analytical clarity and simplicity, they rely on unrealistic assumptions in real commodity markets (constant volatility). The Schwartz one-factor model introduced a more realistic structure by incorporating convenience yields, while the Heston model further advanced realism by modelling volatility as a stochastic, mean-reverting process correlated with price.

Through numerical implementation and comparative analysis, we showed that all 3 models produce similar option valuations under short tenors and stable volatility. However, only models with stochastic or path-dependent volatility can fully capture real-world phenomena such as volatility smiles, term structures and dynamic risk.

Bridging theory with application, we proposed a long call option strategy on LME aluminium futures expiring in February 2026. Supported by macroeconomic analysis, volatility diagnostics and Monte Carlo simulation, the trade appears attractively priced. It benefits from a favourable risk-reward profile: limited downside (premium paid), meaningful convexity above the break-even point and measurable mispricing relative to both theoretical value and skew-adjusted implied volatility curves.

Our findings suggest that while classical models remain useful for pricing and risk management, practitioners should complement them with volatility surface analysis and macroeconomic insight to identify mispricing and structure asymmetric trades. In markets like aluminium, where structural drivers and supply constraints coexist with low implied volatility, such tools offer a powerful framework for informed decision-making.

Closing thought of the ESFS Quant Team: sometimes the trade isn't in the price… it's in the mispricing.

In options markets, price and volatility are two sides of the same coin. Every option premium reflects a view on direction, timing and uncertainty. But while many focus on forecasting where the market is going, real opportunities often lie in understanding where the market is mispricing risk.

This report showed that even a simple call option, when undervalued by models and volatility surfaces, can offer asymmetric returns. The insight is this: the smartest trades aren't always about being right. They're about paying less when others overpay for safety or underpay for risk.

So, the next time you find a quiet corner of the volatility surface, ask yourself: is this mispricing telling me something everyone else is missing?

# Appendix & Bibliography

**References**

Chen, J. (2015, April 10). *How to circumvent the limitations of the Black-Scholes model*. Investopedia. https://www.investopedia.com/articles/active-trading/041015/how-circumvent-limitations-blackscholes-model.asp

Çetin, U., Jarrow, R., Protter, P., & Warachka, M. (2005). *Pricing options in an extended Black-Scholes economy with illiquidity: Theory and empirical evidence*. *Review of Financial Studies*.

Black, F., & Scholes, M. (1973). *The pricing of options and corporate liabilities. Journal of Political Economy, 81*(3), 637–654.

Merton, R. C. (1973). *Theory of rational option pricing. The Bell Journal of Economics and Management Science, 4*(1), 141–183.

Schwartz, E. S. (1997). *The stochastic behaviour of commodity prices: Implications for valuation and hedging. The Journal of Finance, 52*(3), 923–973.

Heston, S. L. (1993). *A closed-form solution for options with stochastic volatility with applications to bond and currency options. The Review of Financial Studies, 6*(2), 327–343.

Longstaff, F. A., & Schwartz, E. S. (2001). *Valuing American options by simulation: A simple least-squares approach. The Review of Financial Studies, 14*(1), 113–147.

Bloomberg Terminal. (2025). *Aluminium market data and option pricing metrics*. Bloomberg L.P.

Financial Times. (2025). *Aluminium market analysis: supply constraints and price outlook*.

Financial Times. (2025). *Aluminium market tightens amid Chinese production caps*.

Financial Times. (2025). *LME aluminium stocks fall as supply disruptions persist*.