

MINERVA INVESTMENT MANAGEMENT SOCIETY

Quantitative Research Division



Option Pricing via Fourier Methods

A Survey of Theory, Methods, Strengths and Limitations

Ianina Tarca

Head of Quantitative Research

Afonso Tavares da Silva

Quantitative Research Analyst

Academic Year 2025/2026

Contents

1	Introduction to Fourier Transforms	3
1.1	Theoretical Introduction	3
1.2	The Risk-Neutral Pricing Connection	4
2	The Motivation for Fourier Methods in Option Pricing	4
2.1	The Black-Scholes Baseline and Its Failures	4
2.2	Richer Models and Their Characteristic Functions	5
2.3	Computational Speed and Model Calibration	5
3	Survey of Fourier Pricing Methods	5
3.1	Heston's Direct Inversion (1993)	6
3.2	Carr & Madan (1999): The FFT Breakthrough	6
3.3	Lewis (2001): A Cleaner Formulation	6
3.4	Chourdakis (2004): The Fractional FFT	7
3.5	Lord et al. (2008): The CONV Method	7
3.6	Fang & Oosterlee (2008): The COS Method	7
4	Fourier Methods in Perspective: Strengths and Limitations	7
4.1	Strengths	7
4.2	Limitations	8
4.3	Overall Assessment	8

Abstract

This paper surveys Fourier-based methods for pricing European options under models with known characteristic functions. After introducing the Fourier transform and the role of characteristic functions in probability, we explain how risk-neutral option valuation can be reformulated in the frequency domain. We then motivate the approach by discussing the limitations of the Black–Scholes model and the tractability of richer models such as Heston, Merton, Variance Gamma, and Normal Inverse Gaussian. The main Fourier pricing methods reviewed are Heston’s direct inversion, the Carr–Madan FFT, Lewis’s contour-integral approach, the fractional FFT, the CONV method, and the COS method. We conclude by comparing the strengths and limitations of Fourier methods, emphasizing their speed, modularity, and effectiveness in vanilla option calibration.

1 Introduction to Fourier Transforms

1.1 Theoretical Introduction

The Fourier transform is best understood as a change of perspective on functions. A function $f(x)$ defined in the spatial (or time) domain can be equivalently expressed in terms of its *frequency content*—how it decomposes into oscillatory components. Specifically, f can be represented as a superposition of complex exponentials,

$$f(x) = \int_{-\infty}^{\infty} \hat{f}(\xi) e^{i\xi x} d\xi,$$

where $\hat{f}(\xi)$ measures the contribution of the oscillation at frequency ξ . The Fourier transform itself is defined by

$$\hat{f}(\xi) = \int_{-\infty}^{\infty} f(x) e^{-i\xi x} dx,$$

which can be read as projecting f onto each oscillatory mode.

The utility of this representation lies in how it simplifies operations. Convolution in the spatial domain,

$$(f * g)(x) = \int_{-\infty}^{\infty} f(y) g(x - y) dy,$$

becomes pointwise multiplication in the frequency domain,

$$\widehat{(f * g)}(\xi) = \hat{f}(\xi) \hat{g}(\xi),$$

and differentiation becomes multiplication by $i\xi$:

$$\widehat{f'}(\xi) = i\xi \hat{f}(\xi).$$

The frequency domain thus turns complex functional operations into simple algebra, which is often the transform's decisive advantage in both analysis and computation.

This perspective extends naturally to probability distributions. A probability density can be viewed as a function, and its Fourier transform provides a frequency-domain representation of randomness. This leads to the central object of interest in this setting: the characteristic function.

Let X be a real-valued random variable. Its *characteristic function* is defined as

$$\phi_X(u) = \mathbb{E}[e^{iuX}], \quad u \in \mathbb{R}.$$

When X admits a density f_X , this is simply the Fourier transform of f_X . Moments, when they exist, are recovered by differentiation at the origin:

$$\mathbb{E}[X^n] = \frac{1}{i^n} \frac{d^n}{du^n} \phi_X(u) \Big|_{u=0}.$$

Crucially, ϕ_X always exists, even when the density has no closed form, making it a robust representation of randomness.

More importantly, probabilities can be recovered directly from ϕ_X without ever knowing the density. The cumulative distribution function $F_X(x) = \mathbb{P}(X \leq x)$ is obtained via the *Gil-Pelaez inversion theorem* [1]:

$$F_X(x) = \frac{1}{2} - \frac{1}{\pi} \int_0^\infty \Im \left(\frac{e^{-iux} \phi_X(u)}{u} \right) du,$$

where $\Im(\cdot)$ denotes the imaginary part. This is the key bridge to option pricing: since option values are expectations under a probability measure, and expectations can be computed directly from ϕ_X , pricing can be carried out entirely in the frequency domain, without requiring an explicit density.

1.2 The Risk-Neutral Pricing Connection

Under the risk-neutral measure $\mathbb{Q}$, the value of a European call with maturity T and strike K is

$$C_0 = e^{-rT} \mathbb{E}^{\mathbb{Q}}[(S_T - K)^+].$$

Working with the log-price $X_T = \log S_T$, this expectation is an integral against the risk-neutral distribution of X_T , which is generally unavailable in closed form.

If the characteristic function $\phi_{X_T}(u)$ is known, however, the distribution can be recovered via Gil-Pelaez inversion, and the expectation computed without requiring the density. Fourier methods are applied to the *log-price* rather than the price level for good reason: log-returns are additive over time, their characteristic functions factorize under independent increments, and many standard models yield exponentially affine forms that are analytically convenient.

The central principle is therefore: *if the characteristic function of the log-price is known, option prices can be computed.* Fourier methods exploit this by reducing pricing to the evaluation of integrals in the frequency domain.

2 The Motivation for Fourier Methods in Option Pricing

2.1 The Black-Scholes Baseline and Its Failures

Under the Black-Scholes model [2], the log-price is Gaussian, with

$$X_T = \log S_T \sim \mathcal{N}\left(\log S_0 + \left(r - \frac{1}{2}\sigma^2\right)T, \sigma^2 T\right),$$

and characteristic function

$$\phi_{X_T}(u) = \exp\left(iu\left(\log S_0 + \left(r - \frac{1}{2}\sigma^2\right)T\right) - \frac{1}{2}\sigma^2 u^2 T\right).$$

The density is known explicitly, and the Gil-Pelaez inversion reduces to the familiar closed-form price. Numerical Fourier methods are therefore unnecessary here.

The importance of Black-Scholes lies not in what it requires analytically, but in what it fails to capture empirically. The Gaussian assumption implies thin tails and symmetry, producing a flat implied volatility surface. Market data consistently contradict this: the

volatility smile (the systematic variation of implied volatility across strikes and maturities) reflects fat tails, negative skew, and jump risk. This is a structural failure, not a calibration issue, and motivates models with richer dynamics: precisely where the Fourier approach becomes indispensable.

2.2 Richer Models and Their Characteristic Functions

Models designed to overcome Black-Scholes follow a common pattern: they introduce more realistic dynamics, lose closed-form densities, but retain tractable characteristic functions.

Merton (1976) [3]. Adding Poisson jumps to the diffusion produces a return distribution that is an infinite mixture of Gaussians. The density is not available in closed form, but the characteristic function remains explicit.

Heston (1993) [4]. Stochastic volatility driven by a CIR process generates skew and smile effects endogenously. The resulting characteristic function is exponentially affine and admits a closed-form expression.

Variance Gamma (Madan, Carr & Chang, 1998). A pure-jump Lévy process with high flexibility for fitting implied volatility surfaces. The density is complicated; the characteristic function is simple and explicit.

Normal Inverse Gaussian (Barndorff-Nielsen, 1997). Another Lévy process with semi-heavy tails and closed-form characteristic function, well-suited to capturing empirical return distributions.

Affine jump-diffusions (Duffie, Pan & Singleton, 2000). A general framework encompassing many of the above models; by construction, the characteristic function is always available in closed form.

The pattern is consistent: while densities become analytically intractable, characteristic functions remain accessible. Fourier inversion is therefore not merely convenient—it is often the only analytical route to option pricing in these models.

2.3 Computational Speed and Model Calibration

Calibrating models such as Heston or Merton to the volatility surface means choosing parameters so that model prices match observed market prices:

$$\min_{\theta} \sum_i w_i \left(C^{\text{model}}(K_i, T_i; \theta) - C^{\text{market}}(K_i, T_i) \right)^2.$$

This requires pricing hundreds of options per iteration across thousands of optimizer steps, imposing strict computational demands. Monte Carlo methods are too slow; PDE solvers do not scale efficiently across strikes.

Fourier methods, by contrast, price an entire strip of strikes at a given maturity in a single computation, reducing pricing to integral evaluations typically completed in milliseconds. This makes them indispensable for real-time calibration in practice.

3 Survey of Fourier Pricing Methods

The mathematical foundation of Section 1 and the financial motivation of Section 2 converge on a practical question: given that the characteristic function is known, how

do you turn it into option prices efficiently? The answer has evolved over three decades through a sequence of methods, each addressing a concrete limitation of its predecessor. The common thread is the same throughout: work in frequency space where the model is tractable, invert only as much as necessary, and exploit whatever computational structure the pricing integral admits.

3.1 Heston's Direct Inversion (1993)

The first practical use of Fourier methods in option pricing is Heston's direct inversion formula [4]. The call price is expressed in terms of two risk-neutral probabilities,

$$C_0 = S_0 P_1 - K e^{-rT} P_2,$$

where P_1 and P_2 are obtained from the characteristic function of the log-price via Gil-Pelaez inversion [1]:

$$P_j = \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \Re \left(\frac{e^{-iu \log K} \phi_{X_T}^{(j)}(u)}{iu} \right) du, \quad j = 1, 2,$$

where $\phi_{X_T}^{(j)}$ denotes a modified characteristic function depending on the payoff. This approach is numerically stable and accurate, and was the first widely used implementation of Fourier pricing. Its limitation is that it computes prices one strike at a time, making it too slow for full surface calibration.

3.2 Carr & Madan (1999): The FFT Breakthrough

Carr and Madan [5] provide the key computational advance by reformulating option pricing to exploit the Fast Fourier Transform. The central idea is to introduce a dampening factor $\alpha > 0$ so that the modified call price

$$\tilde{C}(k) = e^{\alpha k} C(k), \quad k = \log K,$$

is square-integrable and admits a Fourier transform. This reveals a convolution structure in frequency space, allowing an entire strip of strikes at a fixed maturity to be computed simultaneously in a single $O(N \log N)$ pass. The main limitation is that the discretization couples frequency spacing to strike spacing, restricting flexibility when pricing at arbitrary or irregularly spaced strikes.

3.3 Lewis (2001): A Cleaner Formulation

Lewis [6] reformulates option pricing as a single Fourier integral directly in terms of the characteristic function. By shifting the integration contour into the complex plane, the payoff becomes integrable without introducing an explicit dampening parameter. For a European call, the price is

$$C_0 = S_0 - \frac{\sqrt{S_0 K} e^{-rT}}{\pi} \int_0^\infty \Re \left(\frac{e^{-iu \log K} \phi_{X_T}(u - \frac{i}{2})}{u^2 + \frac{1}{4}} \right) du.$$

This formulation is more compact and numerically stable than Heston's original inversion, while avoiding ad hoc modifications of the payoff. Like direct inversion, however, it prices one strike at a time.

3.4 Chourdakis (2004): The Fractional FFT

Chourdakis [7] extends the Carr–Madan framework by introducing the fractional Fast Fourier Transform. The key improvement is decoupling the frequency and strike grids, removing the rigid linkage imposed by the standard FFT. Option prices can then be computed directly at arbitrary strike locations without interpolation, while retaining the same $O(N \log N)$ complexity. This makes the method particularly suited for calibration, where market strikes are typically irregularly spaced.

3.5 Lord et al. (2008): The CONV Method

Lord et al. [8] extend Fourier pricing to early-exercise products by formulating option valuation as a sequence of convolutions. The continuation value at each exercise date is obtained by convolving the value function with the transition density of the log-price; in the Fourier domain, these convolutions reduce to pointwise multiplications evaluated efficiently via the FFT. The result is the CONV method, which preserves the $O(N \log N)$ complexity of Carr–Madan while accommodating Bermudan and American-style features.

3.6 Fang & Oosterlee (2008): The COS Method

Fang and Oosterlee [9] propose a different approach to Fourier inversion by expanding the density in a cosine series over a truncated domain. For $X_T = \log S_T$ on $[a, b]$, the density is approximated as

$$f_{X_T}(x) \approx \sum_{k=0}^{N-1} A_k \cos\left(\frac{k\pi(x-a)}{b-a}\right),$$

where the coefficients A_k are explicit functions of ϕ_{X_T} . This leads to a pricing formula involving a finite cosine series, with no numerical integration required. The method achieves exponential convergence, reaching machine precision with far fewer terms than FFT-based approaches. The COS method is widely regarded as one of the most accurate and efficient techniques available for European option pricing.

4 Fourier Methods in Perspective: Strengths and Limitations

4.1 Strengths

Model agnosticity. The pricing engine is entirely decoupled from the model layer. Swapping Heston for Variance Gamma or NIG requires only substituting the characteristic function; the FFT or COS algorithm is untouched. This modularity is a decisive practical advantage in production environments where multiple models must coexist.

Speed and scalability. A full strip of strikes at one maturity requires a single FFT call. Since maturities are independent, the entire volatility surface can be priced in milliseconds—a requirement that Monte Carlo cannot meet and that makes real-time calibration feasible.

Analytical Greeks. Delta, Gamma, and Vega are obtained by differentiating under the Fourier integral analytically, avoiding finite-difference approximations and their associated noise and overhead.

Payoff flexibility. Within the COS framework, different payoffs correspond to different coefficient vectors computed analytically once. The same inversion machinery prices vanilla calls, digitals, and asset-or-nothing options without modification.

4.2 Limitations

Branch-cut instability. The characteristic function of models such as Heston involves complex logarithms whose branch cuts must be handled carefully. Naive implementations produce incorrect prices, particularly at long maturities.

Parameter sensitivity. The dampening parameter α in Carr–Madan and the truncation interval $[a, b]$ in the COS method depend on the model and its parameters. Poor choices produce numerical instability or truncation error that can be difficult to detect without careful diagnostics.

European vanillas by default. The basic framework naturally prices European options. Barriers, Americans, and Bermudans require non-trivial extensions, such as the CONV method [8], which add significant implementation complexity.

Closed-form characteristic function required. The approach depends entirely on the characteristic function being available analytically. For rough volatility models or models specified via neural SDEs, direct application is not possible without approximation.

Curse of dimensionality. Fourier inversion does not extend naturally to high dimensions. Multi-asset and basket options remain outside the practical reach of these methods.

4.3 Overall Assessment

Against Monte Carlo, Fourier methods win decisively on speed for vanilla calibration but lose on generality. Against PDE solvers, they are easier to deploy across models but less natural for American options and high-dimensional problems. The asset class for which they perform best—European vanillas under models with known characteristic functions, calibrated to a full implied volatility surface—is precisely where they are universally adopted in practice, and where no competing method offers a comparable combination of speed, accuracy, and model flexibility.

Fourier methods provide a powerful framework for pricing options whenever the characteristic function of the log-price is known in closed form. Their importance stems from a simple but deep idea: while probability densities are often analytically inaccessible in realistic asset-pricing models, characteristic functions remain tractable. This makes frequency-domain pricing not only mathematically elegant but computationally indispensable. Among the methods surveyed, FFT-based approaches are particularly valuable for calibration, while the COS method stands out for its accuracy and efficiency in European vanilla pricing. Despite their limitations in high dimensions and in products with path dependence or early exercise, Fourier methods remain one of the central tools of modern quantitative finance.

References

- [1] Gil-Pelaez, J. (1951). *Note on the inversion theorem*. *Biometrika*, 38(3–4), 481–482.

- [2] Black, F. & Scholes, M. (1973). *The pricing of options and corporate liabilities*. Journal of Political Economy, 81(3), 637–654.
- [3] Merton, R.C. (1976). *Option pricing when underlying stock returns are discontinuous*. Journal of Financial Economics, 3(1–2), 125–144.
- [4] Heston, S.L. (1993). *A closed-form solution for options with stochastic volatility with applications to bond and currency options*. Review of Financial Studies, 6(2), 327–343.
- [5] Carr, P. & Madan, D. (1999). *Option valuation using the fast Fourier transform*. Journal of Computational Finance, 2(4), 61–73.
- [6] Lewis, A.L. (2001). *A simple option formula for general jump-diffusion and other exponential Lévy processes*. SSRN Working Paper. Available at: <https://ssrn.com/abstract=282110>
- [7] Chourdakis, K. (2004). *Option pricing using the fractional FFT*. Journal of Computational Finance, 8(2), 1–18.
- [8] Lord, R., Fang, F., Bervoets, F. & Oosterlee, C.W. (2008). *A fast and accurate FFT-based method for pricing early-exercise options under Lévy processes*. SIAM Journal on Scientific Computing, 30(4), 1678–1705.
- [9] Fang, F. & Oosterlee, C.W. (2008). *A novel pricing method for European options based on Fourier-cosine series expansions*. SIAM Journal on Scientific Computing, 31(2), 826–848.