Quantitative Finance and Economics
https://www.aimspress.com/journal/QFE

QFE, 10(3): 405–437.
DOI: 10.3934/QFE.2026017
Received: 22 December 2025
Revised: 03 April 2026
Accepted: 14 May 2026
Published: 14 July 2026

Research article

Option pricing: A K–theoretic perspective on diffusion models

Ioannis P. ZOIS[1,2,*]

[1] Exeter College, Turl Street, OX1 3DP, Oxford, UK
[2] Public Power Corporation (PPC) S.A., 30 Chalkokondyli Street, GR 104 32, Athens, Greece

* **Correspondence:** Email: i.zois@exeter.oxon.org, i.zois@ppcgroup.com.

Abstract: **Background:** Continuous–time option pricing under diffusion models leads, after the usual risk–neutral reduction and standard changes of variables, to parabolic partial differential equations generated by second–order elliptic operators with Gaussian heat kernels. This suggests a geometric reformulation of diffusion pricing in terms of Dirac–type operators and index theory. The purpose of this article was to develop such a structural viewpoint and to identify the invariant information it could reveal about payoff spaces.
Methods: We studied stylised Bachelier and Black–Scholes–Merton diffusion generators on compactified state spaces and realised them as Laplace–type operators of the form $-D^2$ for suitable Dirac–type operators. On the two–torus equipped with a $U(1)$ line bundle of degree k, we analysed the associated twisted Dirac operator using heat–kernel methods, the Atiyah–Singer index theorem, and an explicit theta–function construction of zero modes. We then related the resulting invariant sector to a finite–market diagnostic based on payoff vectors and Gram matrices.
Results: We obtained a Dirac realisation of stylised pricing generators on S^1 and T^2, proved that the twisted Dirac operator on T^2 has index k, and gave an explicit Jacobi theta–function basis for its zero modes. These zero modes spanned a finite–dimensional diffusion–invariant subspace of payoff space in the compactified model. In addition, the finite–market example showed how the same viewpoint led to a rank/kernel diagnostic for payoff redundancy and constrained market dimension.
Conclusions: Index–theoretic methods provide a coherent structural language for diffusion–based option pricing. Although the compactified constructions are stylised rather than literal market models, they isolate stable geometric features of pricing generators and motivate finite–dimensional diagnostics that can complement standard quantitative–finance workflows.

Keywords: option pricing; Bachelier model; Black–Scholes–Merton model; Dirac operators; Atiyah–Singer index theorem; spectral triples; K–theory; theta functions

JEL Codes: G13, C02, C65

1. Introduction

The standard arbitrage–free paradigm for continuous–time financial markets asserts that under a risk–neutral measure, derivative prices are given by discounted conditional expectations. Analytically, this leads to parabolic PDEs whose generators are second–order elliptic operators. In the simplest cases, such as Bachelier's additive Brownian motion (Bachelier, 1900), Einstein's Brownian motion and diffusion (Einstein, 1905), and the Black–Scholes–Merton model (Black and Scholes, 1973; Merton, 1973; Björk, 2009; Shreve, 2004), the transition density of the underlying process is governed by a heat equation and the pricing semigroup is a heat semigroup. From a modern finance perspective, these classical models belong to a broader quantitative landscape including volatility modelling, calibration, and comparison between normal and lognormal frameworks (Gatheral, 2006; Schachermayer and Teichmann, 2008; Choi et al., 2022; Alòs and García-Lorite, 2025).

Recent work has also explored unified diffusion frameworks that interpolate between Bachelier and Black–Scholes–Merton type settings, further underscoring the continuing value of structural comparisons between normal and lognormal pricing regimes (Lindquist et al., 2024).

Independently, heat kernels and second–order elliptic operators are central in global analysis and index theory. The McKean–Singer formula expresses the index of an elliptic Dirac–type operator as a supertrace of the associated heat semigroup, and the Atiyah–Singer index theorem identifies this index with a topological invariant of the underlying manifold and vector bundle (Atiyah and Singer, 1963; Berline et al., 1992; Gilkey, 1984; Bismut, 1986; Lawson and Michelsohn, 1989; Freed, 2021; Roe, 1998; Reutter, 2015). Noncommutative geometry further generalises this framework through spectral triples and K–homology/K–theory pairings (Connes, 1994; Zois, 2000, 2010). Thus, the same analytic objects that appear naturally in diffusion pricing—second–order operators, heat semigroups, and heat kernels—also appear in the index–theoretic study of Dirac operators.

The present article exploited this parallel structure. Its aim was not to replace the probabilistic approach to option pricing, nor to propose a new calibration framework, but rather to isolate and study a structural layer of diffusion–based pricing models. More precisely, we showed how stylised pricing generators could be realised as squares of Dirac–type operators on compact state spaces, and how index theory and theta–function techniques could then be used to identify finite–dimensional invariant sectors of payoff space. Economically, the point is not that these sectors immediately deliver an implementable trading rule; rather, they provide a disciplined form of dimensional reduction inside the stylised diffusion semigroup. The zero–mode sector acts as a spectrally distinguished component of payoff space, while the complementary modes are smoothed and damped. This was why the later finite–market Gram–matrix analysis should be read as a diagnostic shadow of the geometric construction: it detects redundancy, constrained effective payoff-space dimension, and near-degenerate directions among traded payoffs.

A central theme throughout the article was the passage from the usual stochastic/PDE description to an operator–theoretic one:

$$\text{diffusion dynamics} \;\longrightarrow\; \text{pricing generator} \;\longrightarrow\; \text{heat semigroup} \;\longrightarrow\; \text{Dirac–type square and index.}$$

This chain did not by itself produce empirical calibration or a trading strategy in the present article, but it revealed invariant information that is not transparent at the level of the pricing PDE alone. In particular, once the diffusion generator was embedded into a Dirac–type framework, the kernel of the corresponding operator and its index became natural objects to study.

For readers coming primarily from quantitative finance, four objects should be kept in view. First, a Dirac–type operator D is a first–order operator whose square $-D^2$ reproduces the elliptic core of the pricing generator in the compactified models. Second, the twisting line bundle $L_k \to T^2$ carries the topological datum $k = c_1(L_k)[T^2]$, which controls the index. Third, the zero modes are the elements of $\ker D_A$; these are exactly the payoff configurations left fixed by the stylised backward diffusion semigroup. Fourth, K–theory and K–homology provide the stable bookkeeping framework in which vector bundles/projections and elliptic operators can be paired to produce integer invariants. For the purposes of this article, however, the essential content is concrete rather than abstract: the kernel defines a finite–dimensional distinguished sector, and the index measures its stable signed size. For general background on topological and operator-algebraic K–theory, see (Atiyah, 1967; Karoubi, 1978; Wegge-Olsen, 1993; Blackadar, 1998; Higson and Roe, 2000); for the noncommutative-geometric viewpoint relevant here, see (Connes, 1994; Zois, 2000, 2010).

From the finance side, it is important to state the scope carefully. The compactified models studied below should be regarded as stylised geometric laboratories rather than literal market models. Passing from the real line or plane to S^1 or T^2 is a mathematical idealisation used to make the operator theory globally tractable and to expose the invariant structures explicitly. Likewise, when we speak of stability or of a topologically stable sector, the meaning is mathematical stability under smooth deformations within the stylised model, not automatic economic robustness under liquidity constraints, transaction costs, stochastic volatility, jumps, or model misspecification. Similarly, the article did not claim to provide a full hedging theory with discrete rebalancing and frictions. Rather, it identified structural payoff sectors and motivated a finite–market diagnostic based on payoff vectors, Gram matrices, and numerical rank/kernel computations. In this respect, the present work was closer to a structural or diagnostic viewpoint than to a feasibility–first implementation framework, and should be read as complementary to more directly practitioner–oriented approaches such as (Halidias, 2025).

A second point requiring clarification is the role of drift and discounting. In the classical Bachelier and Black–Scholes–Merton frameworks these terms appear naturally in the stochastic dynamics and in the pricing PDE. However, after the standard risk–neutral reduction and the usual changes of variables, the central analytic object is the second–order diffusion operator. Since it is this elliptic core that interacts directly with the heat kernel and the square of a Dirac–type operator, the present article isolated that core in order to develop the index–theoretic viewpoint. The zero–drift formulas appearing later should therefore be read as a deliberate structural simplification, not as a claim that realistic markets have zero drift.

Main results. Informally, the three core contributions of the article were:

(i) A Dirac–type realisation of stylised pricing generators on compact state spaces (circle and torus).
(ii) An explicit theta–function description of the zero modes of a twisted Dirac operator on T^2, together with a rigorous index formula.
(iii) A financial interpretation of these zero modes as a finite–dimensional, diffusion–invariant sector of payoff space, together with a finite–market diagnostic based on payoff vectors and Gram matrices.

These were formalised as the following theorem–level statements, proved in Sections 6, 7 and 8 respectively. At an intuitive level, Theorem 1 identifies the stable integer carried by the twisted compactified model, Theorem 2 shows that this invariant is represented concretely by an explicit theta–function basis of zero modes, and Theorem 3 interprets those zero modes as the persistent part of the stylised backward diffusion semigroup.

Theorem 1 (Index of twisted Dirac pricing operator on T^2)**.** *Let T^2 be the flat two–torus with its standard spin structure, and let $L_k \to T^2$ be a complex line bundle of degree $k \in \mathbb{Z}$ equipped with a unitary connection of constant curvature. Let D_A be the associated twisted Dirac operator and $L_A = -D_A^2$ the corresponding Laplace–type pricing generator. Then*

$$\mathrm{index}(D_A) = k,$$

and this integer is stable under any smooth perturbation of the metric and the connection that preserves $c_1(L_k) = k$.

Theorem 2 (Theta–function basis of zero modes)**.** *In the setting of Theorem 1 with $k > 0$, the positive–chirality kernel* $\ker D_A^+$ *has complex dimension k and admits an explicit orthonormal basis*

$$\psi_{+,j}(z,\bar{z}) = \mathcal{N}\, e^{-\pi k z\bar{z}}\, \vartheta\begin{bmatrix} \frac{j}{k} \\ 0 \end{bmatrix}(kz \mid ki), \qquad j = 0, 1, \ldots, k-1,$$

where $z = x + iy$ is the complex coordinate on the square torus $T^2 = \mathbb{C}/(\mathbb{Z} + i\mathbb{Z})$, ϑ is the Jacobi theta function with characteristics, and $\mathcal{N}$ is a normalisation constant. For $k < 0$, an analogous statement holds for $\ker D_A^-$.

Theorem 3 (Diffusion–invariant payoff sector)**.** *Let $L_A = -D_A^2$ be as in Theorem 1, and consider the backward pricing PDE*

$$\frac{\partial \Psi}{\partial \tau} = L_A \Psi, \qquad \tau \geq 0,$$

for spinor–valued payoffs $\Psi(\tau, \cdot) \in L^2(T^2, S \otimes L_k)$. Then:

(a) For any initial condition $\Psi_0 \in \ker D_A$ the unique mild solution is $\Psi(\tau) = \Psi_0$ for all $\tau \geq 0$; that is, Ψ_0 is diffusion–invariant.
(b) The space of diffusion–invariant payoffs has complex dimension $|\mathrm{index}(D_A)| = |k|$ and is spanned by the theta–function zero modes described in Theorem 2.

The article is organised as follows. Section 2 revisits diffusion–based option pricing and the heat equation, with emphasis on Bachelier, Einstein and Black–Scholes–Merton, and clarifies how drift and discounting are handled when one isolates the diffusion–driven elliptic core. Section 3 discusses the dissipative features of the heat flow and their financial interpretation in simple models. Section 4 reviews the heat–kernel proof of the Atiyah–Singer index theorem in a form tailored to our purposes. Sections 5 and 6 construct Dirac–type operators on S^1 and T^2 whose squares reproduce pricing generators on compactified state spaces, and prove Theorem 1. Section 7 gave an explicit description of the zero modes of this twisted Dirac operator via classical theta functions and proves Theorem 2. Section 8 interpreted these theta–function zero modes as diffusion–invariant payoff structures and proves Theorem 3, including a concrete numerical toy example. Section 9 formulated the spectral triple (A, H, D_A) and outlines how K–theoretic index pairings can be used to define structural invariants of payoff spaces. Section 10 contained an explicit finite–market example and a practitioner–facing Gram–matrix diagnostic for redundancy and constrained market dimension. Section 11 concluded with a discussion of limitations, possible extensions, and applications.

2. Option pricing and the heat equation

This section recalled the diffusion models that motivated the article and made explicit the analytic point that will be used later. In both the additive (Bachelier) and multiplicative (Black–Scholes–Merton) settings, the pricing problem is governed by a second–order parabolic operator. After the standard risk–neutral reduction, discounting, and elementary changes of variables, one isolates a diffusion–driven elliptic core whose semigroup is a heat semigroup. It is this structural core, rather than the full market model with all lower–order terms, that is carried forward into the Dirac–type framework developed in the later sections.

2.1. Bachelier's additive Brownian motion

In his 1900 thesis *Théorie de la spéculation* (Bachelier, 1900), Bachelier modelled an asset price S_t as an additive Brownian motion with drift,

$$S_t = S_0 + \mu t + \sigma W_t, \tag{1}$$

where W_t is standard Brownian motion, μ is the drift, and σ is the volatility. Then S_T is normally distributed, and a European call option with strike K and maturity T has present value

$$C_0 = e^{-rT}\,\mathbb{E}[(S_T - K)^+], \tag{2}$$

which can be computed explicitly from the Gaussian distribution.

In this model, the probability density $p(s,t)$ of S_t satisfies the Fokker–Planck equation

$$\frac{\partial p}{\partial t} = -\mu\frac{\partial p}{\partial s} + \frac{\sigma^2}{2}\frac{\partial^2 p}{\partial s^2}. \tag{3}$$

Thus the drift contributes only a first–order transport term, while the diffusion contributes the second–order elliptic term. If one introduces the translated variable

$$\xi = s - \mu t$$

and writes

$$q(\xi,t) = p(\xi + \mu t, t),$$

then q satisfies the pure heat equation

$$\frac{\partial q}{\partial t} = \frac{\sigma^2}{2}\frac{\partial^2 q}{\partial \xi^2}. \tag{4}$$

So even when $\mu \neq 0$, the underlying analytic mechanism is still the heat equation after passage to a moving frame. This is the first reason why the zero–drift formulas used later are mathematically natural: they isolate the diffusion core without denying the possible presence of drift in realistic financial modelling.

From a pricing point of view, one may similarly regard the backward pricing operator in the additive model as consisting of a second–order diffusion term plus lower–order contributions coming from drift and discounting. After the usual discounting and a simple change of variables, the same heat operator reappears. In this sense, the present article did not rely on a literal claim that additive models have zero drift; rather, it uses the fact that the second–order part is the structurally relevant one for the heat–kernel and Dirac–square viewpoint.

Although classical, the Bachelier model remains relevant in modern quantitative finance as a normal model and as a comparison point for lognormal modelling. The relation between Bachelier and Black–Scholes pricing has been analysed carefully in (Schachermayer and Teichmann, 2008), while practical use of the Bachelier framework and volatility conversion issues are discussed in (Choi et al., 2022). Recent work continues to study Bachelier implied volatility and its refinements; see, for example, (Alòs and García-Lorite, 2025).

2.2. Einstein's Brownian motion, diffusion and calibration

Einstein's 1905 paper on Brownian motion (Einstein, 1905) placed this picture on a firm physical and probabilistic foundation. Starting from the molecular–kinetic theory of heat, he derived the diffusion equation for the probability density of a suspended particle undergoing random motion, and identified the diffusion coefficient D in terms of microscopic parameters such as temperature, viscosity and particle size.

On the one hand, considering a one–dimensional random walk with independent increments, Einstein obtained the relation

$$\langle x^2(t) \rangle = 2Dt, \tag{5}$$

so that the probability density $p(x, t)$ of the particle's position satisfies

$$\frac{\partial p}{\partial t} = D \frac{\partial^2 p}{\partial x^2}, \tag{6}$$

with D the diffusion constant. This is precisely the heat equation, with Gaussian fundamental solution and variance $2Dt$.

On the other hand, by combining the kinetic picture with Stokes' law for the viscous drag on a spherical particle of radius r in a fluid of viscosity η, Einstein derived the celebrated Stokes–Einstein relation

$$D = \frac{k_B T}{6\pi\eta r}, \tag{7}$$

where k_B is Boltzmann's constant and T the absolute temperature. This formula connects the macroscopic diffusion coefficient D to microscopic parameters of the system.

From the point of view of stochastic differential equations, a one–dimensional Brownian motion with diffusion coefficient D can be written as

$$dX_t = \sqrt{2D}\, dW_t, \tag{8}$$

and the density of X_t solves $\partial_t p = D\partial_x^2 p$. In a standard financial diffusion model for the (log–)price,

$$dX_t = \sigma\, dW_t, \tag{9}$$

the corresponding density satisfies

$$\frac{\partial p}{\partial t} = \frac{\sigma^2}{2} \frac{\partial^2 p}{\partial x^2}. \tag{10}$$

Thus, at the PDE level, the diffusion constant in Einstein's setting is identified with

$$D = \frac{\sigma^2}{2}. \tag{11}$$

In physical Brownian motion, (7) expresses D in terms of (k_B, T, η, r) and thereby provides a bridge between microscopic dynamics and macroscopic diffusion. In financial markets, we do not have a universally accepted microscopic model of order flow and liquidity of the same status as molecular kinetics; instead, the effective diffusion coefficient $D = \sigma^2/2$ is calibrated from market data. Conceptually, however, the role of D is the same: it measures the strength of random fluctuations and sets the time scale of dispersion of the probability distribution. The Einstein picture thus provides a useful template for thinking about the interpretation of σ (or D) as an emergent large–scale parameter encoding microscopic randomness, whether physical or financial.

2.3. The Black–Scholes–Merton equation

Modern mathematical finance replaced additive Brownian motion by geometric Brownian motion (GBM). In the risk–neutral measure, the asset price S_t satisfies the stochastic differential equation (Black and Scholes, 1973; Merton, 1973; Björk, 2009; Shreve, 2004)

$$dS_t = rS_t\,dt + \sigma S_t\,dW_t, \tag{12}$$

where r is the risk–free rate and σ the volatility. By Itô's lemma, $\log S_t$ is Gaussian and S_t is lognormal.

Let $V(S,t)$ be the value of a derivative claim with maturity T and payoff $\Phi(S_T)$. Under the standard assumptions of the Black–Scholes framework (absence of arbitrage, frictionless trading, continuous hedging), the delta–hedging argument shows that V satisfies the Black–Scholes–Merton PDE

$$\frac{\partial V}{\partial t} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 V}{\partial S^2} + rS\frac{\partial V}{\partial S} - rV = 0, \tag{13}$$

with terminal condition $V(T,S) = \Phi(S)$.

Introduce the time–to–maturity variable $\tau = T - t$ and the log–moneyness variable

$$x = \log(S/K),$$

where $K > 0$ is a reference strike. A standard conjugation of the form

$$V(S,t) = e^{\alpha x + \beta\tau} u(x,\tau), \tag{14}$$

with suitable constants α, β, removes the first–order and zeroth–order terms and transforms (13) into the heat equation

$$\frac{\partial u}{\partial \tau} = \frac{1}{2}\sigma^2 \frac{\partial^2 u}{\partial x^2}, \tag{15}$$

for $u(x,\tau)$, with initial condition $u(\cdot,0)$ determined by the payoff. Solving this heat equation with the Gaussian heat kernel and transforming back yields the closed–form Black–Scholes formula for a European call.

This calculation was important for the present article because it showed that the full Black–Scholes–Merton pricing operator is conjugate to a pure diffusion operator. In other words, after discounting and a standard change of unknown, the essential analytic object is again the second–order elliptic operator

$$\frac{\sigma^2}{2}\frac{\partial^2}{\partial x^2}.$$

It is this elliptic core that later admits a natural realisation as the negative square of a Dirac–type operator on a compactified state space.

To summarise, both the additive and multiplicative diffusion models lead, after standard reductions, to heat equations or operators conjugate to heat equations. This is the only aspect of the pricing models used in the sequel. The later geometric constructions do not aim to reproduce every market feature; they isolate the common heat–kernel structure that underlies the diffusion pricing semigroup.

3. Heat equation, dissipation and Bachelier's model

The purpose of this section is to isolate the qualitative effect of the heat semigroup that later reappears in the Dirac–type framework. The key point is simple: diffusion smooths. In the forward equation for probability densities, this smoothing spreads mass and reduces concentration; in the backward pricing equation, written in time–to–maturity, it regularises payoff profiles. This is the analytic mechanism behind the later appearance of finite–dimensional invariant sectors: most modes are smoothed or damped, while exceptional modes may remain fixed.

Consider first the classical one–dimensional heat equation

$$\frac{\partial u}{\partial t} = \kappa \frac{\partial^2 u}{\partial x^2}, \qquad \kappa > 0, \tag{16}$$

with suitable decay or periodic boundary conditions. If u is interpreted as a density, then the total mass

$$H(t) = \int u(x,t)\, dx \tag{17}$$

is conserved. At the same time, the heat flow smooths spatial oscillations. For example, the Dirichlet energy

$$E(t) = \int \left|\partial_x u(x,t)\right|^2 dx \tag{18}$$

is nonincreasing in t, so sharp gradients are progressively washed out. Thus the heat semigroup preserves mass while dissipating high–frequency structure.

This picture is already present in Bachelier's model (Bachelier, 1900). In the zero–drift case, the density $p(s,t)$ of the asset price satisfies

$$\frac{\partial p}{\partial t} = \frac{\sigma^2}{2} \frac{\partial^2 p}{\partial s^2}. \tag{19}$$

Hence the transition density is Gaussian, total probability is conserved, and the distribution becomes more spread out as time increases. A useful way to express this loss of concentration is through the Fisher information

$$I(t) = \int \frac{(p_s)^2}{p}\, ds, \tag{20}$$

which decreases along the heat flow when the solution is sufficiently regular. The density therefore becomes less peaked and less sensitive to small spatial variations.

Einstein's derivation of the diffusion equation from microscopic random motion (Einstein, 1905) gives this smoothing a clear physical meaning. The cumulative effect of many small random kicks drives the system away from sharp, localised initial profiles and toward more diffuse ones. In a stylised financial setting the same PDE mechanism is present: a diffusion model washes out fine structure in the state variable and replaces it by Gaussian averaging.

For option pricing, however, one must distinguish carefully between the *forward* heat equation for densities and the *backward* pricing equation for derivative values. In a Bachelier world with zero drift, if $\tau = T - t$ denotes time to maturity and f is the payoff at maturity, then the option value $V(s, \tau)$ satisfies

$$\frac{\partial V}{\partial \tau} = \frac{\sigma^2}{2} \frac{\partial^2 V}{\partial s^2}, \qquad V(s, 0) = f(s). \tag{21}$$

Thus, as a function of τ, the price is obtained by applying the heat semigroup to the payoff:

$$V(\cdot, \tau) = e^{\tau(\sigma^2/2)\partial_s^2} f. \tag{22}$$

Equivalently, V is the Gaussian convolution of f with variance $\sigma^2\tau$.

This observation has an immediate financial interpretation. Nonsmooth payoffs are regularised by diffusion. For example, the kink of a vanilla call payoff at the strike is smoothed once one moves away from maturity. More generally, the pricing semigroup suppresses highly oscillatory components of a payoff profile. This does *not* mean that economic information is destroyed in any naive sense; rather, it means that diffusion pricing averages local irregularities in the state variable. Analytically, the semigroup acts as a smoothing operator.

This smoothing perspective will be important later. On compact state spaces such as S^1 and T^2, the heat semigroup admits a discrete spectral decomposition. Modes corresponding to positive eigenvalues decay, while zero modes remain unchanged. The later Dirac–type construction makes this distinction precise: the invariant payoff sector arises from the kernel of the relevant operator, whereas the remaining modes are governed by the dissipative effect of the heat flow.

To summarise: in physical heat flow, in Einstein's diffusion picture, and in Bachelier–type pricing, the same second–order elliptic mechanism is at work. It preserves total mass in the forward density equation, smooths and spreads profiles, and in the backward pricing problem acts by Gaussian regularisation of payoffs. This common analytic structure is what makes the later passage to heat kernels and Dirac squares natural.

4. Heat kernels and the Atiyah–Singer index theorem

This section recalled, in a form adapted to the present article, the heat–kernel approach to index theory. The reason for including it was simple. In the pricing problems discussed earlier, the central analytic objects were second–order elliptic operators and their heat semigroups. The same objects occur in the study of Dirac–type operators on manifolds. The purpose of the present section was therefore not to reproduce the full theory, but to record the specific facts that would later be used when pricing generators are realised as negative squares of Dirac–type operators.

Let M be a compact Riemannian manifold and let $E \to M$ be a complex vector bundle. A Dirac–type operator

$$D : C^\infty(M, E) \longrightarrow C^\infty(M, E)$$

is a first–order elliptic differential operator whose square D^2 is a second–order elliptic operator of Laplace type. After completion in $L^2(M, E)$, the operator D is essentially self–adjoint, and the heat semigroup

$$e^{-tD^2}, \qquad t > 0,$$

is a smoothing operator with a smooth heat kernel $K_t(x, y)$.

Analytically, this is exactly the same semigroup structure already encountered in diffusion pricing. In the financial setting one starts from a stochastic model, passes to the infinitesimal pricing generator, and then to the associated heat semigroup. In the geometric setting one starts from a Dirac–type operator, passes to its square, and again arrives at a heat semigroup. The bridge used in this article was therefore:

$$\text{pricing generator } L \quad \rightsquigarrow \quad L = -D^2 \quad \rightsquigarrow \quad e^{\tau L} = e^{-\tau D^2}.$$

The point is not that every pricing model is literally geometric, but that once such a realisation exists, the heat–kernel machinery of index theory becomes available.

The heat kernel admits a short–time asymptotic expansion along the diagonal,

$$K_t(x, x) \sim \frac{1}{(4\pi t)^{n/2}} \sum_{j=0}^{\infty} t^j a_j(x), \qquad t \downarrow 0, \tag{23}$$

where $n = \dim M$ and the coefficients $a_j(x)$ are locally computable in terms of the metric, curvature, and the connection on the twisting bundle (Gilkey, 1984; Bismut, 1986; Berline et al., 1992; Reutter, 2015; Roe, 1998). Thus the short–time behaviour of the heat kernel encodes geometric information.

When M is even–dimensional and the bundle is $\mathbb{Z}_2$–graded, one writes

$$D = \begin{pmatrix} 0 & D^- \\ D^+ & 0 \end{pmatrix},$$

so that

$$D^2 = \begin{pmatrix} D^-D^+ & 0 \\ 0 & D^+D^- \end{pmatrix}.$$

The analytic index is then defined by

$$\operatorname{index}(D) = \dim\ker D^+ - \dim\ker D^-. \tag{24}$$

A basic fact is that this integer is stable under suitable perturbations: it depends only on the elliptic homotopy class of the operator, not on fine analytic details.

The McKean–Singer formula expresses the index as a supertrace of the heat semigroup (Atiyah and Singer, 1963; Freed, 2021):

$$\operatorname{index}(D) = \operatorname{Tr}(e^{-tD^-D^+}) - \operatorname{Tr}(e^{-tD^+D^-}) = \operatorname{Str}(e^{-tD^2}), \qquad t > 0. \tag{25}$$

The crucial point is that the right–hand side is independent of t. For large t, the heat semigroup suppresses all positive modes and the supertrace sees only the zero modes. For small t, the heat kernel asymptotics turn the same quantity into an integral of local curvature expressions. Equating these two descriptions yields the index theorem.

For a twisted spin Dirac operator D_E on a compact spin manifold, the Atiyah–Singer formula takes the form

$$\operatorname{index}(D_E) = \int_M \widehat{A}(TM) \wedge \operatorname{ch}(E), \tag{26}$$

where $\widehat{A}(TM)$ is the $\widehat{A}$–genus of the tangent bundle and $\mathrm{ch}(E)$ is the Chern character of the twisting bundle (Atiyah and Singer, 1963; Lawson and Michelsohn, 1989; Berline et al., 1992; Freed, 2021). In dimension two, this simplifies considerably, and for the twisted torus model studied later the index reduces to the first Chern number of the line bundle.

This is the key fact used in the sequel. Once a pricing generator is realised as $-D^2$ for a suitable twisted Dirac operator, the kernel of D becomes a distinguished finite–dimensional sector, and its dimension is controlled by a topological quantity. In the compactified two–asset model studied later, this distinguished sector will be interpreted as a diffusion–invariant payoff sector. The financial interpretation should be read carefully: what is proved is a mathematical statement about the kernel of a stylised geometric operator, not a complete implementable hedging theorem.

For completeness, let us also recall the role of heat kernels in the semigroup picture. If $\{\phi_j\}$ is an orthonormal eigenbasis for D^2 with

$$D^2\phi_j = \lambda_j\phi_j, \qquad \lambda_j \geq 0,$$

then

$$e^{-tD^2}\phi_j = e^{-t\lambda_j}\phi_j. \tag{27}$$

Hence strictly positive modes decay exponentially as $t \to \infty$, while zero modes remain fixed. This is the spectral form of the smoothing principle already seen in Section 3, and it is the analytic reason why kernels of Dirac–type operators naturally define invariant sectors.

To summarise, the heat–kernel proof of the Atiyah–Singer theorem relies on the same analytic ingredients that appear in diffusion pricing: elliptic operators, heat semigroups, spectral damping, and short–time kernel asymptotics. This is why the passage from pricing generators to Dirac squares is mathematically natural and why index theory becomes relevant in the compactified models constructed below.

5. A one–dimensional toy model: Black–Scholes on a circle

Before passing to the two–dimensional torus, it is useful to record the simplest compactified example. Consider a one–asset model in which the log–price variable is compactified to a circle

$$S^1 = \mathbb{R}/2\pi\mathbb{Z},$$

with angular coordinate $\theta \in [0, 2\pi)$. As stressed earlier, this compactification is a stylised mathematical idealisation whose purpose is to make the operator theory global and explicit; it is not intended as a literal market model.

After the standard risk–neutral reduction and the elimination of lower–order terms discussed in Section 2, the diffusion core of the one–dimensional pricing operator is

$$L = \frac{\sigma^2}{2}\frac{d^2}{d\theta^2}. \tag{28}$$

We showed that this operator is the negative square of a natural Dirac–type operator on S^1.

The standard Dirac operator on the circle is

$$D_0 = -i\frac{d}{d\theta}, \tag{29}$$

acting on complex–valued functions (equivalently, on one–dimensional spinors). Rescaling by the volatility gives

$$D = \sqrt{\frac{\sigma^2}{2}}\, D_0 = -i\sqrt{\frac{\sigma^2}{2}}\,\frac{d}{d\theta}. \tag{30}$$

Therefore

$$D^2 = -\frac{\sigma^2}{2}\frac{d^2}{d\theta^2}, \tag{31}$$

and hence

$$L = -D^2. \tag{32}$$

Thus the pricing semigroup acting backward in time–to–maturity τ is

$$e^{\tau L} = e^{-\tau D^2}, \tag{33}$$

which is exactly the heat semigroup of the Dirac square. This is the simplest instance of the analytic bridge used throughout the article.

The spectral picture is completely explicit. The Fourier modes

$$\phi_n(\theta) = e^{in\theta}, \qquad n \in \mathbb{Z}, \tag{34}$$

satisfy

$$D_0\phi_n = n\phi_n, \qquad D^2\phi_n = \frac{\sigma^2}{2}n^2\phi_n, \qquad L\phi_n = -\frac{\sigma^2}{2}n^2\phi_n. \tag{35}$$

Hence

$$e^{\tau L}\phi_n = e^{-(\sigma^2/2)n^2\tau}\phi_n. \tag{36}$$

All nonzero Fourier modes are exponentially damped as τ increases, while the constant mode ϕ_0 remains fixed. This is the compact one–dimensional version of the smoothing principle discussed in Section 3.

The circle example therefore already contains the essential analytic mechanism: a pricing generator is realised as the negative square of a Dirac operator, and the corresponding pricing semigroup is a heat semigroup. What it does *not* yet provide is a nontrivial even index. Since S^1 is odd–dimensional, the ordinary circle Dirac operator does not lead in this untwisted setting to the nonzero integer index phenomenon that appears in even dimensions. The one–dimensional model is therefore only a toy bridge. To obtain a genuinely nontrivial topological invariant, one must pass to the two–dimensional torus and introduce a twist by a line bundle of nonzero Chern class.

6. Two assets on a torus and a nonzero index

We passed from the one–dimensional toy model to the first genuinely topological case. Consider a stylised two–asset setting with state variables given by compactified log–prices. To keep the notation consistent with the theta–function construction of the next section, we worked on the square torus

$$T^2 \;=\; \mathbb{R}^2/\mathbb{Z}^2 \;\cong\; \mathbb{C}/(\mathbb{Z} + i\mathbb{Z}),$$

with coordinates $(x, y) \in [0, 1)^2$ and flat metric

$$ds^2 = dx^2 + dy^2.$$

As in the circle example, this compactification is a stylised mathematical idealisation whose purpose is to make the operator theory global and explicit. It should not be read as a literal market model. The point is to expose, in the simplest possible compact setting, the relation between diffusion generators, Dirac squares, and topological index.

6.1. *Flat Dirac operator and diffusion generator*

Let $S \to T^2$ be the spinor bundle. Since the torus is flat and parallelisable, S is trivial of rank 2, and we may use the Pauli matrices

$$\sigma_1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \qquad \sigma_2 = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}. \tag{37}$$

The flat Dirac operator is

$$D_0 = -i(\sigma_1 \partial_x + \sigma_2 \partial_y), \tag{38}$$

acting on spinor fields $\psi : T^2 \to \mathbb{C}^2$. Using the anticommutation relations of the Pauli matrices, one checks immediately that

$$D_0^2 = -\Delta\, I_2, \qquad \Delta = \partial_x^2 + \partial_y^2. \tag{39}$$

Rescaling by the volatility parameter, define

$$D = \sqrt{\frac{\sigma^2}{2}}\, D_0. \tag{40}$$

Then

$$D^2 = -\frac{\sigma^2}{2}\, \Delta\, I_2. \tag{41}$$

Accordingly, the diffusion core of the two–dimensional pricing generator is

$$L_0 = \frac{\sigma^2}{2}\, \Delta, \tag{42}$$

so that

$$L_0 = -D^2. \tag{43}$$

Thus, exactly as in the one–dimensional case,

$$e^{\tau L_0} = e^{-\tau D^2}.$$

This is the compactified two–asset analogue of the heat–equation reduction discussed in Section 2. For simplicity we wrote the operator with equal volatilities and without mixed second derivatives. The purpose is to isolate the common elliptic core, not to model the full covariance structure of a realistic multi–asset market.

6.2. *Twisting by a line bundle of degree k*

The untwisted operator on T^2 still has trivial index. To obtain a nonzero topological invariant, one must twist by a nontrivial complex line bundle. Let

$$L_k \longrightarrow T^2$$

be a Hermitian line bundle of degree $k \in \mathbb{Z}$, equipped with a unitary connection A whose curvature F_A satisfies

$$\frac{i}{2\pi} \int_{T^2} F_A = k. \tag{44}$$

Equivalently, $c_1(L_k)[T^2] = k$.

Let ∇^A denote the covariant derivative on L_k, and let $S \otimes L_k$ be the twisted spinor bundle. The twisted Dirac operator is defined by

$$D_A = -i\sqrt{\frac{\sigma^2}{2}}(\sigma_1 \nabla^A_x + \sigma_2 \nabla^A_y). \tag{45}$$

Its square is a Laplace–type operator. More precisely, one has the standard Weitzenböck identity

$$D_A^2 = -\frac{\sigma^2}{2}\,\Delta_A\, I_2 + \frac{\sigma^2}{2}\, c(F_A), \tag{46}$$

where Δ_A is the covariant Laplacian and $c(F_A)$ denotes Clifford multiplication by the curvature 2–form. Since T^2 is flat, there is no scalar curvature term.

Equation (46) is the precise version of the heuristic used earlier: the dominant second–order part is again a diffusion operator, while the twist contributes a zeroth–order curvature term. Analytically, this is the two–dimensional compact analogue of a diffusion generator with an external field. Financially, however, the twist should be regarded as part of the stylised geometric construction, not as a literal market force.

6.3. *Index computation*

We computed the index of D_A. Since T^2 is an even–dimensional compact spin manifold, the twisted Dirac operator splits into chiral parts

$$D_A = \begin{pmatrix} 0 & D_A^- \\ D_A^+ & 0 \end{pmatrix},$$

and its index is

$$\operatorname{index}(D_A) = \dim\ker D_A^+ - \dim\ker D_A^-. \tag{47}$$

By the Atiyah–Singer index theorem for twisted Dirac operators (Atiyah and Singer, 1963; Lawson and Michelsohn, 1989; Freed, 2021),

$$\operatorname{index}(D_A) = \int_{T^2} \widehat{A}(TT^2) \wedge \operatorname{ch}(L_k). \tag{48}$$

Because the torus is flat, one has

$$\widehat{A}(TT^2) = 1.$$

Since L_k is a line bundle,

$$\mathrm{ch}(L_k) = 1 + c_1(L_k),$$

and only the degree–two component contributes to the integral. Therefore

$$\mathrm{index}(D_A) = \int_{T^2} c_1(L_k) = k. \tag{49}$$

This is exactly the statement of Theorem 1.

The conclusion is important for the sequel. The twisted compactified pricing operator has a distinguished finite–dimensional zero–mode sector whose signed dimension is controlled by the topological degree k of the twisting bundle. For $k > 0$, the positive–chirality kernel has dimension k and the negative one is trivial; for $k < 0$, the roles are reversed. The explicit construction of these zero modes, in terms of theta functions, is the content of the next section.

Thus the torus model is the first place where the heat–kernel/Dirac viewpoint produces a genuinely nontrivial integer invariant. This is the mathematical core behind the later interpretation of a diffusion–invariant payoff sector.

7. Explicit theta–function construction of zero modes

In this section we constructed explicitly the zero modes of the twisted Dirac operator D_A on the square torus T^2 with a $U(1)$ connection of constant curvature, in a holomorphic gauge adapted to the complex structure. The aim is to make Theorem 2 completely explicit. For $k > 0$, the positive–chirality zero modes are described by classical Jacobi theta functions; for $k < 0$, the analogous statement holds for negative chirality.

7.1. Complex structure and holomorphic gauge

We identified the square torus with the elliptic curve

$$T^2 \;\cong\; \mathbb{C}/(\mathbb{Z} + i\mathbb{Z}),$$

and write

$$z = x + iy, \qquad x, y \in \mathbb{R}, \qquad (x, y) \sim (x + 1, y) \sim (x, y + 1).$$

Thus the complex structure parameter is

$$\tau = i,$$

and the flat metric is

$$ds^2 = dx^2 + dy^2 = dz\, d\bar{z}.$$

Let $L_k \to T^2$ be the degree–k line bundle introduced in Section 6. In a local holomorphic gauge, the twisted Chern connection may be written so that its $(0, 1)$ component is

$$\nabla^A_{\bar{z}} = \frac{\partial}{\partial \bar{z}} + \pi k z. \tag{50}$$

This local formula realises a constant–curvature connection representing $c_1(L_k) = k$. Globally, sections of L_k are not ordinary periodic functions on $\mathbb{C}$, but obey quasi–periodicity conditions determined by the line bundle. In the present gauge, these conditions are encoded naturally by theta functions on the elliptic curve; see, for example, (Farkas and Kra, 1992; Griffiths and Harris, 1978).

7.2. Dirac operator and the zero–mode equation

On a Riemann surface, the chiral pieces of the twisted Dirac operator are, up to a harmless constant factor, the twisted Cauchy–Riemann operators. More precisely, for the operator constructed in Section 6, one has

$$D_A = \begin{pmatrix} 0 & D_A^- \\ D_A^+ & 0 \end{pmatrix}, \qquad D_A^+ \sim \sqrt{\frac{\sigma^2}{2}}\,\bar{\partial}_A,$$

where

$$\bar{\partial}_A := \nabla^A_{\bar{z}}.$$

Since the scalar factor $\sqrt{\sigma^2/2}$ does not affect the kernel, the positive–chirality zero modes are exactly the solutions of

$$D_A^+\psi_+ = 0 \quad \Longleftrightarrow \quad \bar{\partial}_A\psi_+ = 0 \quad \Longleftrightarrow \quad \left(\frac{\partial}{\partial\bar{z}} + \pi kz\right)\psi_+(z,\bar{z}) = 0. \tag{51}$$

Equation (51) has the standard local factorisation

$$\psi_+(z,\bar{z}) = e^{-\pi kz\bar{z}}\, f(z), \tag{52}$$

where f is holomorphic. Indeed, substituting

$$\psi_+(z,\bar{z}) = e^{-\pi kz\bar{z}} f(z,\bar{z})$$

into (51) and using

$$\frac{\partial}{\partial\bar{z}}(e^{-\pi kz\bar{z}}) = -\pi kz\, e^{-\pi kz\bar{z}},$$

one finds

$$\bar{\partial}_A\psi_+ = e^{-\pi kz\bar{z}}\,\frac{\partial f}{\partial\bar{z}}.$$

Hence $\bar{\partial}_A\psi_+ = 0$ if and only if f is holomorphic.

The global problem is therefore reduced to determining which holomorphic functions f are compatible with the quasi–periodicity of the line bundle L_k. This is exactly the classical theta–function problem on an elliptic curve (Farkas and Kra, 1992; Griffiths and Harris, 1978).

7.3. Theta functions and holomorphic sections of L_k

For $\tau = i$, recall the Jacobi theta function with characteristics $a, b \in \mathbb{R}$:

$$\vartheta\begin{bmatrix} a \\ b \end{bmatrix}(z \mid \tau) := \sum_{n\in\mathbb{Z}} \exp\Big(\pi i(n+a)^2\tau + 2\pi i(n+a)(z+b)\Big). \tag{53}$$

Classically, holomorphic sections of a degree–k line bundle on $\mathbb{C}/(\mathbb{Z}+i\mathbb{Z})$ are represented by theta functions of level k; see again (Farkas and Kra, 1992; Griffiths and Harris, 1978).

For $k > 0$, define

$$f_j(z) := \vartheta\begin{bmatrix} \frac{j}{k} \\ 0 \end{bmatrix}(kz \mid ki), \qquad j = 0, 1, \ldots, k-1. \tag{54}$$

Each f_j is holomorphic and has exactly the quasi–periodicity needed to descend to a holomorphic section of L_k. Moreover, the space of holomorphic sections $H^0(T^2, L_k)$ has dimension k, and the family $\{f_j\}_{j=0}^{k-1}$ is a basis.

Substituting (54) into (52), we obtained the explicit zero modes

$$\psi_{+,j}(z,\bar{z}) = \mathcal{N}\, e^{-\pi k z\bar{z}}\, \vartheta\begin{bmatrix} \frac{j}{k} \\ 0 \end{bmatrix}(kz \mid ki), \qquad j = 0, 1, \ldots, k-1, \tag{55}$$

where $\mathcal{N}$ is a normalisation constant chosen so that the resulting spinors are orthonormal in $L^2(T^2, S \otimes L_k)$.

It follows immediately that:

- each $\psi_{+,j}$ solves $\bar{\partial}_A \psi_{+,j} = 0$;
- the family $\{\psi_{+,j}\}_{j=0}^{k-1}$ is linearly independent;
- every solution of $\bar{\partial}_A \psi_+ = 0$ is a unique linear combination of the $\psi_{+,j}$.

Therefore

$$\ker D_A^+ \cong \ker \bar{\partial}_A \cong \operatorname{span}\{\psi_{+,0}, \ldots, \psi_{+,k-1}\}, \qquad \dim\ker D_A^+ = k \qquad (k > 0).$$

For $k < 0$, the same argument applies after reversing chirality: then $\ker D_A^+ = \{0\}$ and $\dim\ker D_A^- = |k|$.

This proves Theorem 2.

7.4. Fourier–series viewpoint

It is useful to connect the theta–function basis with a more elementary Fourier description. On the real torus $\mathbb{R}^2/\mathbb{Z}^2$, ordinary periodic functions decompose in the Fourier basis

$$e_{n,m}(x,y) = e^{2\pi i(nx+my)}, \qquad n, m \in \mathbb{Z}. \tag{56}$$

A section of L_k is not strictly periodic, but in the present gauge it can be written locally as a Gaussian factor multiplied by a Fourier series whose coefficients are constrained by the bundle quasi–periodicity.

Writing

$$\psi_+(x,y) = e^{-\pi k(x^2+y^2)} \sum_{n\in\mathbb{Z}} c_n \exp(2\pi i n y + 2\pi k n x), \tag{57}$$

the quasi–periodicity relations force the coefficients to satisfy

$$c_{n+k} = c_n$$

up to the corresponding gauge phase. Thus only k independent coefficient patterns remain. The resulting k admissible Fourier superpositions are exactly the theta–function combinations in (55).

From this point of view, each zero mode is a Gaussian–modulated Fourier superposition singled out by the compatibility of the holomorphic condition with the line–bundle topology.

7.5. Summary of the explicit basis

For the square torus

$$T^2 = \mathbb{C}/(\mathbb{Z} + i\mathbb{Z})$$

and a degree–k line bundle L_k with constant–curvature connection, the positive–chirality zero modes for $k > 0$ are given explicitly by

$$\psi_{+,j}(z,\bar{z}) = \mathcal{N}\, e^{-\pi k z\bar{z}}\, \vartheta\begin{bmatrix} \frac{j}{k} \\ 0 \end{bmatrix}(kz \mid ki), \qquad j = 0, 1, \ldots, k-1. \tag{58}$$

These spinors form a basis of $\ker D_A^+$ and realise concretely the index formula

$$\operatorname{index}(D_A) = \dim\ker D_A^+ - \dim\ker D_A^- = k.$$

The financial interpretation of these zero modes is deferred to Section 8.

8. Financial interpretation of theta–function zero modes

We interpreted the theta–function zero modes constructed in Section 7 in financial terms and prove Theorem 3. The interpretation must be read carefully. The point of this section was not to inflate the financial claim, but to state precisely what the zero modes mean inside the stylised compactified model. What is established here is a structural statement about the backward diffusion semigroup: the kernel of the twisted Dirac operator defines a finite–dimensional subspace of payoff configurations that is left fixed by the stylised diffusion dynamics once the usual discounting has been separated off, while the complementary directions are smoothed and damped. In that sense, the zero–mode sector provides a disciplined form of spectral filtering and dimensional reduction in payoff space. This was a statement about the geometry and spectrum of the pricing operator, not a complete trading or hedging theorem in a market with discrete rebalancing, transaction costs, stochastic volatility, jumps, or other frictions.

8.1. State space, payoffs, and the backward pricing semigroup

In the two–asset compactified model, the state space is the square torus

$$T^2 = \mathbb{C}/(\mathbb{Z} + i\mathbb{Z}),$$

with coordinate

$$z = x + iy, \qquad (x, y) \in [0, 1)^2.$$

Scalar payoffs at maturity are represented by bounded functions on T^2, while in the twisted setting developed in Sections 6 and 7 the natural objects are spinor–valued sections

$$\Psi \in L^2(T^2, S \otimes L_k).$$

The diffusion core of the pricing dynamics is governed by the Laplace–type operator

$$L_A = -D_A^2,$$

where D_A is the twisted Dirac operator associated with the degree–k line bundle $L_k \to T^2$. Ignoring discounting factors that can be separated off by the standard transformations discussed in Section 2, the backward pricing equation takes the form

$$\frac{\partial \Psi}{\partial \tau} = L_A \Psi = -D_A^2 \Psi, \qquad \tau \geq 0, \tag{59}$$

with $\tau = T - t$ the time to maturity. Its solution is given by the heat semigroup

$$\Psi(\tau) = e^{\tau L_A}\Psi_0 = e^{-\tau D_A^2}\Psi_0. \tag{60}$$

Thus the financial interpretation of the zero modes is governed by a simple spectral fact: the semigroup damps positive modes of D_A^2, whereas it leaves the kernel unchanged.

8.2. *Proof of the diffusion–invariant payoff theorem*

We proved Theorem 3.

Suppose first that $\Psi_0 \in \ker D_A$. Then

$$D_A\Psi_0 = 0, \qquad D_A^2\Psi_0 = 0, \qquad L_A\Psi_0 = 0.$$

Therefore the unique mild solution of (59) with initial condition Ψ_0 is

$$\Psi(\tau) = e^{\tau L_A}\Psi_0 = \Psi_0, \qquad \tau \geq 0.$$

So every zero mode is fixed by the backward diffusion semigroup.

Conversely, suppose that a payoff configuration Ψ_0 satisfies

$$e^{\tau L_A}\Psi_0 = \Psi_0 \qquad \text{for all } \tau \geq 0.$$

Differentiating at $\tau = 0$ in the generator sense yields

$$L_A\Psi_0 = 0,$$

hence

$$-D_A^2\Psi_0 = 0.$$

Since D_A is self–adjoint on the compact manifold T^2, we have

$$\langle \Psi_0, D_A^2\Psi_0\rangle = \|D_A\Psi_0\|^2.$$

Thus $D_A^2\Psi_0 = 0$ implies

$$\|D_A\Psi_0\|^2 = 0,$$

and therefore

$$D_A\Psi_0 = 0.$$

Hence the fixed points of the semigroup are exactly the zero modes of D_A.

This proves part (a) of Theorem 3.

For part (b), Section 7 gave an explicit basis of theta–function zero modes. If $k > 0$, then

$$\dim\ker D_A^+ = k, \qquad \ker D_A^- = \{0\},$$

and if $k < 0$, the roles of the two chiralities are reversed, with

$$\dim\ker D_A^- = |k|, \qquad \ker D_A^+ = \{0\}.$$

Therefore in all cases

$$\dim\ker D_A = |k| = |\operatorname{index}(D_A)|,$$

and this space is spanned by the explicit theta–function modes constructed in Section 7. This proves part (b) of Theorem 3.

8.3. Interpretation as a structural payoff–space decomposition

The theorem above gives a precise and useful interpretation of the zero modes. Define

$$\mathcal{Z}_k := \ker D_A \subset L^2(T^2, S \otimes L_k).$$

Then $\mathcal{Z}_k$ is a finite–dimensional subspace of payoff configurations that is invariant under the backward diffusion semigroup:

$$e^{\tau L_A}\mathcal{Z}_k = \mathcal{Z}_k, \qquad e^{\tau L_A}\Psi = \Psi \text{ for all } \Psi \in \mathcal{Z}_k.$$

Let

$$P_0 : L^2(T^2, S \otimes L_k) \to \mathcal{Z}_k$$

denote the orthogonal projection onto the zero–mode sector. Then any payoff configuration Ψ_0 admits the orthogonal decomposition

$$\Psi_0 = P_0\Psi_0 + (I - P_0)\Psi_0. \tag{61}$$

Under the semigroup, the first term is preserved and the second is damped:

$$e^{\tau L_A}\Psi_0 = P_0\Psi_0 + e^{\tau L_A}(I - P_0)\Psi_0.$$

Since the nonzero spectrum of D_A^2 is strictly positive on the orthogonal complement of $\ker D_A$, one has exponential decay there. Writing $\lambda_1 > 0$ for the smallest strictly positive eigenvalue of D_A^2, we obtained

$$\left\| e^{\tau L_A}(I - P_0)\Psi_0 \right\| \le e^{-\lambda_1 \tau} \|(I - P_0)\Psi_0\|. \tag{62}$$

So the backward diffusion dynamics separates an arbitrary payoff into:

(i) a finite–dimensional invariant component $P_0\Psi_0$, and
(ii) a transient component $(I - P_0)\Psi_0$ that is smoothed and damped.

In economic terms, the decomposition (61) is best read as a stylised dimensional–reduction mechanism. The projection $P_0\Psi_0$ extracts the spectrally persistent component of the payoff, while the complementary part represents transient directions filtered out by the diffusion semigroup. This does not by itself produce an executable trading strategy, but it does motivate diagnostics of effective payoff-space dimension, redundancy, and constrained directions. That is precisely why the finite-market Gram–matrix analysis of Section 10 is included: it is the finite-dimensional shadow of the same structural idea.

8.4. On Greeks, hedging language, and implementability

In classical quantitative finance, sensitivities of prices with respect to model inputs and state variables are encoded by the Greeks; see, for instance, (Björk, 2009; Shreve, 2004; Gatheral, 2006). In the present setting, the zero–mode condition

$$D_A\Psi = 0$$

imposes first–order differential constraints on the components of the payoff configuration. In that sense, the directional sensitivities with respect to the two compactified log–price variables are not independent.

However, it is important to state clearly what this does and does not mean. The present article did *not* prove that every zero mode corresponds to a self–financing trading strategy built from a specified set of traded instruments under discrete rebalancing, bid–ask spreads, margin constraints, liquidity frictions, stochastic volatility, or jump risk. Those are separate feasibility questions. What the article showed is that, within the stylised compactified model, the zero modes form a finite–dimensional payoff sector singled out by the geometry and spectrum of the pricing operator.

For this reason, expressions such as "hedge structure" should be read here in a deliberately limited sense: the zero modes describe payoff configurations annihilated by the stylised diffusion operator after discounting is removed. They are best understood as structural or diagnostic objects, and as motivation for the finite-dimensional redundancy tests developed later in the article, rather than as a complete implementation recipe for trading practice.

8.5. Mathematical stability versus economic robustness

The index computation of Section 6 implies that the dimension of $\ker D_A$ depends only on the Chern number k of the twisting bundle. In particular, if the geometric data are deformed smoothly while preserving the topological class $c_1(L_k) = k$, then the integer

$$\dim\ker D_A = |k|$$

remains unchanged. This is the precise mathematical sense in which the zero–mode sector is stable.

This stability should not be confused with economic robustness in a practical market sense. A trading or risk–management procedure may be sensitive to sampling frequency, transaction costs, quote noise, liquidity constraints, stochastic volatility, or jump effects even when a mathematical index remains unchanged. Accordingly, the word "protected" in the present context refers only to stability of the kernel dimension inside the stylised geometric model, not to an automatic guarantee of market performance. These structural features may nevertheless persist, at least in part, in more realistic models, but that is a matter for explicit analysis and practical verification rather than something established here.

8.6. A numerical toy example: projection onto the zero–mode sector

We described a simple numerical toy example in the case $k = 1$. Its purpose is purely illustrative: to show how an arbitrary payoff may be decomposed into a zero–mode component and a transient component as in (61). It is not intended as a real–market calibration exercise.

Model. Let

$$T^2 = \mathbb{R}^2/\mathbb{Z}^2, \qquad (x, y) \in [0, 1)^2,$$

and take the degree–one line bundle $L_1 \to T^2$ with the constant–curvature connection used in Sections 6 and 7. Then $\dim\ker D_A = 1$, and we denoted by ψ_0 the normalised zero mode:

$$\psi_0(z, \bar{z}) = \mathcal{N}\, e^{-\pi z\bar{z}}\, \vartheta\begin{bmatrix}0\\0\end{bmatrix}(z \mid i), \tag{63}$$

with $\|\psi_0\|_{L^2} = 1$.

Payoff. Consider the scalar payoff

$$f(x, y) = \max\bigl(\cos(2\pi x) + \cos(2\pi y) - 1,\ 0\bigr), \tag{64}$$

and embed it into the first spinor component:

$$\Psi_f(x,y) = \begin{pmatrix} f(x,y) \\ 0 \end{pmatrix}.$$

Projection coefficient. The orthogonal projection onto $\ker D_A$ is

$$P_0\Psi = \langle \psi_0, \Psi \rangle\, \psi_0.$$

Hence the zero–mode part of the payoff is

$$\Psi_f^{(0)} = P_0\Psi_f = c_f\,\psi_0, \qquad c_f := \langle \psi_0, \Psi_f \rangle. \tag{65}$$

The transient part is

$$\Psi_f^{\perp} = \Psi_f - \Psi_f^{(0)}.$$

Discrete approximation. Choose a uniform $N \times N$ grid

$$x_p = \frac{p}{N}, \qquad y_q = \frac{q}{N}, \qquad p, q = 0, \ldots, N-1.$$

Then the projection coefficient is approximated by the Riemann sum

$$c_f^{(N)} = \frac{1}{N^2} \sum_{p,q=0}^{N-1} \overline{\psi_{0,1}(x_p, y_q)}\, f(x_p, y_q), \tag{66}$$

where $\psi_{0,1}$ is the first component of ψ_0. The projected payoff on the grid is

$$\Psi_f^{(0,N)} = c_f^{(N)}\psi_0,$$

and the remainder is

$$\Psi_f^{(\perp,N)} = \Psi_f - \Psi_f^{(0,N)}.$$

Backward diffusion. The evolution under the semigroup may then be approximated either by:

(i) a truncated spectral decomposition of the discrete operator approximating L_A, or
(ii) a stable implicit time–stepping scheme for the corresponding discrete heat equation.

In either approach, one expects

$$e^{\tau L_A}\Psi_f \approx \Psi_f^{(0)}$$

for moderate $\tau > 0$, because the zero–mode component is fixed while the orthogonal complement is damped according to (62).

Thus the toy example gives a concrete numerical picture of the theorem: an arbitrary payoff is decomposed into a persistent component in the zero–mode sector and a transient component smoothed out by the diffusion.

8.7. Towards the K–theoretic viewpoint

The explicit theta–function basis makes the zero–mode sector completely concrete. This is important for the next section, where the same data are reinterpreted through the spectral triple (A, H, D_A) in the sense of noncommutative geometry (Connes, 1994). The kernel of D_A is the most visible manifestation of the analytic index, while the spectral triple framework packages the same structure in a way adapted to K–theory and index pairings. In this broader setting, the zero–mode sector plays the role of a distinguished finite–dimensional component singled out by the operator, and the index theorem is the bridge between its analytic and topological descriptions. This is also the natural point of contact with the use of K–theoretic and cyclic methods in related noncommutative settings (Zois, 2000).

9. C*–algebra of payoffs, spectral triples, and K–theory

The preceding sections were formulated at the level of differential operators and explicit zero modes. We recast the same data in the language of noncommutative geometry. The purpose of this section is not to develop the full abstract theory, but to show that the compactified pricing model naturally defines a spectral triple and therefore admits K–theoretic index pairings. This provided a compact algebraic framework for the structural information already identified analytically.

Let

$$A = C(T^2) \tag{67}$$

be the C*–algebra of continuous complex–valued functions on the state space T^2. In the present stylised setting, an element $f \in A$ represents a bounded scalar payoff function depending on the two compactified log–price variables.

Let

$$H = L^2(T^2, S \otimes L_k) \tag{68}$$

be the Hilbert space of square–integrable twisted spinors, and let A act on H by pointwise multiplication:

$$(\pi(f)\Psi)(x, y) = f(x, y)\Psi(x, y). \tag{69}$$

Together with the twisted Dirac operator D_A, this gives the triple

$$(A, H, D_A), \tag{70}$$

which is a spectral triple in the sense of Connes (Connes, 1994); for general background on K–theory, K–homology, and operator algebras, see also (Atiyah, 1967; Karoubi, 1978; Wegge-Olsen, 1993; Blackadar, 1998; Higson and Roe, 2000). The analytic content of this statement is the following:

(i) the algebra A acts faithfully on H;
(ii) the operator D_A is self–adjoint with compact resolvent on the compact manifold T^2;
(iii) for smooth f, the commutator $[D_A, \pi(f)]$ is bounded, since D_A is a first–order differential operator.

Thus the compactified pricing model fits naturally into the standard framework of spectral geometry.

The value of this reformulation is conceptual. In Sections 6–8, the kernel of D_A and its index were obtained by direct differential–geometric analysis. The spectral triple packages the same information in a form adapted to K–homology and index pairings. In particular, the operator D_A determines a K–homology class

$$[D_A] \in K^0(A), \tag{71}$$

while projections in matrix algebras over A define classes in

$$K_0(A). \tag{72}$$

This is standard in analytic K–homology and operator-algebraic K–theory (Higson and Roe, 2000; Blackadar, 1998; Wegge-Olsen, 1993). Concretely, let

$$e \in M_N(A)$$

be a projection, so that $e^2 = e = e^*$. Such a projection represents a K–theory class

$$[e] \in K_0(A).$$

Geometrically, one may think of e as encoding a finite–rank subbundle of the trivial bundle $T^2 \times \mathbb{C}^N$. In the present financial interpretation, it is safer to regard e not as a literal trading strategy, but as a stylised family of payoff configurations or constraints varying over state space.

The pairing between the K–homology class of D_A and the K–theory class of e is the integer

$$\langle [D_A], [e] \rangle = \operatorname{index}(eD_Ae : eH \to eH) \in \mathbb{Z}. \tag{73}$$

This is the natural K–theoretic refinement of the analytic index computed earlier. In particular, if one takes the trivial projection $e = 1$, then

$$\langle [D_A], [1] \rangle = \operatorname{index}(D_A) = k. \tag{74}$$

So the integer obtained in Section 6 is the simplest instance of a K–theoretic index pairing.

This point of view is useful because it separates two layers of structure:

(a) the operator D_A, which encodes the geometric diffusion dynamics through its square $-D_A^2$;
(b) the K–theory class $[e]$, which selects a finite–rank family of states or payoff configurations on which one measures the index.

The pairing (73) then produces a stable integer invariant. In this way, the zero–mode sector identified earlier is not an isolated spectral curiosity; it is the visible analytic part of a broader index–theoretic framework.

It is important, however, not to overstate the financial meaning. The integer

$$\langle [D_A], [e] \rangle$$

is a structural invariant of the stylised compactified model. It does not by itself produce an arbitrage bound, a hedging strategy, or a calibration output. Rather, it records topologically stable information about the interaction between the operator encoding the diffusion core and the K–theory class selected by the projection e. In that sense it is best viewed as a structural or diagnostic invariant, in the same general spirit in which index pairings are used in geometric and noncommutative settings (Connes, 1994; Zois, 2000, 2010).

This also clarified the role of the present section within the article. The compactified torus model provides an infinite–dimensional geometric setting in which the index and the zero–mode sector can be computed exactly. Section 10 then turns to a finite market with finitely many traded payoffs, where one can perform an elementary Gram–matrix/rank analysis. That finite calculation is not itself a K–theory computation in the full operator–algebraic sense, but it should be viewed as a finite–dimensional shadow of the same structural idea: the payoff space may contain redundancies, constrained directions, and a distinguished effective dimension. The next section makes this bridge explicit.

10. A finite–market example: static arbitrage and a K–theoretic dimension count

The previous sections treated continuous–time diffusion models and their Dirac realisations on S^1 and T^2. We turned to a deliberately simple finite market in order to extract a concrete diagnostic that can be used in practice. The point of this section is not to claim that the full geometric machinery is needed to perform elementary linear algebra. Rather, it is to show that the structural viewpoint developed earlier has a clear finite–dimensional counterpart. In the compactified diffusion model, the backward semigroup separates payoff space into a persistent sector and transient damped directions. In the finite market considered here, the analogous questions are: what is the effective dimension of the payoff span, which instruments are redundant, and which directions are nearly degenerate? These questions are captured by the Gram matrix and its numerical rank.

In this finite setting, the relevant invariant is the rank of the payoff span. This is the discrete analogue of the distinguished finite-dimensional sectors encountered earlier in the compactified diffusion model. The resulting construction is elementary, numerically tractable, and useful as a diagnostic for redundancy, constrained payoff-space dimension, and near-collinearity among traded payoffs.

10.1. Set–up: a finite one–maturity option market

Fix a maturity $T > 0$ and consider a market with the following traded instruments at time $t = 0$:

- the discounted underlying payoff S_T,
- a family of n European call options

$$C_{K_1}, \dots, C_{K_n},$$

with the same maturity T and strictly increasing strikes

$$K_1 < \cdots < K_n.$$

Let

$$X_0 := S_T, \qquad X_i := C_{K_i}, \quad i = 1, \dots, n,$$

and collect these payoffs into the vector

$$X = (X_0, \dots, X_n).$$

To make the discussion fully explicit, choose a finite scenario set

$$\Omega = \{\omega_1, \dots, \omega_m\},$$

with associated terminal prices $S_T(\omega_j)$ and strictly positive risk–neutral weights

$$\pi_j > 0, \qquad \sum_{j=1}^{m} \pi_j = 1.$$

In practice, Ω may arise from a recombining tree, a finite-state approximation of a diffusion model, or a Monte Carlo sample with equal or importance weights.

Each traded payoff X_i is represented by the vector

$$x_i \in \mathbb{R}^m, \qquad x_i(\omega_j) = X_i(\omega_j).$$

We endowed the payoff space with the weighted inner product

$$\langle x_i, x_j \rangle := \sum_{k=1}^{m} \pi_k \, x_i(\omega_k) x_j(\omega_k). \tag{75}$$

This is the finite-dimensional analogue of the L^2 inner product used in the earlier continuous models.

10.2. The payoff Gram matrix and its kernel

Define the $(n+1) \times (n+1)$ Gram matrix

$$T_{ij} := \langle x_i, x_j \rangle, \qquad i, j = 0, \dots, n. \tag{76}$$

The matrix T is symmetric and positive semidefinite by construction.

A portfolio

$$\alpha = (\alpha_0, \dots, \alpha_n)^\top \in \mathbb{R}^{n+1}$$

produces the payoff

$$X_\alpha = \sum_{i=0}^{n} \alpha_i X_i,$$

represented by

$$x_\alpha = \sum_{i=0}^{n} \alpha_i x_i.$$

Its squared norm is

$$\|x_\alpha\|^2 = \langle x_\alpha, x_\alpha \rangle = \sum_{i,j=0}^{n} \alpha_i \alpha_j T_{ij} = \alpha^\top T \alpha. \tag{77}$$

Since T is positive semidefinite, it follows that

$$\ker T = \{\alpha \in \mathbb{R}^{n+1} : \|x_\alpha\|^2 = 0\} = \{\alpha : X_\alpha(\omega_j) = 0 \text{ for all } j\}. \tag{78}$$

Thus $\ker T$ consists exactly of the portfolios whose payoff vanishes in every scenario. These are pure redundancies in payoff space.

On the other hand,

$$\operatorname{rank}(T) = \dim \operatorname{span}\{x_0, \dots, x_n\},$$

so the rank gives the effective dimension of the payoff family generated by the traded instruments. If one adds instruments without increasing the rank, then one has added only redundant directions.

10.3. A three–strike toy market

We computed an explicit example. Consider one underlying with maturity T and three equally likely terminal scenarios:

$$S_T(\omega_1) = 80, \qquad S_T(\omega_2) = 100, \qquad S_T(\omega_3) = 120,$$

with

$$\pi_1 = \pi_2 = \pi_3 = \frac{1}{3},$$

and zero interest rates for simplicity.

Take three strikes

$$K_1 = 90, \qquad K_2 = 100, \qquad K_3 = 110,$$

and the four traded instruments

- the underlying $X_0 = S_T$,
- the three calls

$$X_i = C_{K_i} = (S_T - K_i)^+, \qquad i = 1, 2, 3.$$

The corresponding payoff vectors are

$$x_0 = (80, 100, 120),$$

$$x_1 = (0, 10, 30), \qquad x_2 = (0, 0, 20), \qquad x_3 = (0, 0, 10).$$

Using (75) and (76), we obtained

$$T = \frac{1}{3}\begin{pmatrix} 80^2 + 100^2 + 120^2 & 100 \cdot 10 + 120 \cdot 30 & 120 \cdot 20 & 120 \cdot 10 \\ 100 \cdot 10 + 120 \cdot 30 & 10^2 + 30^2 & 30 \cdot 20 & 30 \cdot 10 \\ 120 \cdot 20 & 30 \cdot 20 & 20^2 & 20 \cdot 10 \\ 120 \cdot 10 & 30 \cdot 10 & 20 \cdot 10 & 10^2 \end{pmatrix}$$

Evaluating the entries gives

$$T = \frac{1}{3}\begin{pmatrix} 30800 & 4600 & 2400 & 1200 \\ 4600 & 1000 & 600 & 300 \\ 2400 & 600 & 400 & 200 \\ 1200 & 300 & 200 & 100 \end{pmatrix}. \tag{79}$$

A direct rank computation shows that

$$\operatorname{rank}(T) = 3,$$

hence

$$\dim \ker T = 4 - 3 = 1.$$

There is exactly one linearly independent portfolio with zero payoff in all three scenarios. Solving

$$T\alpha = 0$$

(or equivalently $x_\alpha = 0$), one finds

$$\alpha = (0, 0, 1, -2)^\top, \tag{80}$$

up to a nonzero scalar multiple. This corresponds to the static portfolio

$$X_\alpha = X_2 - 2X_3.$$

Direct inspection of the payoff vectors confirms that

$$x_2 - 2x_3 = (0, 0, 20) - 2(0, 0, 10) = (0, 0, 0),$$

so

$$X_\alpha(\omega_j) = 0, \qquad j = 1, 2, 3.$$

Therefore $X_2 - 2X_3$ is a pure redundancy in this toy market.

This example still illustrates the general point of the section: the Gram matrix detects redundant payoff directions and its kernel identifies the corresponding zero-payoff portfolios. In the present toy market, the redundancy comes from the fact that the calls with strikes 100 and 110 are exactly collinear across the chosen three scenarios.

10.4. Practical interpretation

The preceding example illustrates the general method. Given a finite family of traded payoffs:

1. discretise the pricing model into scenarios Ω with weights π_j;
2. form the payoff vectors x_i;
3. build the Gram matrix T;
4. compute its singular values and numerical rank.

The output has direct practical meaning:

- the rank gives the effective static payoff dimension generated by the chosen instruments;
- the kernel identifies exact model redundancies;
- very small singular values indicate near-redundancies and therefore possible instability, quote-noise sensitivity, or near-arbitrage effects.

In applications one does *not* use exact rank in floating-point arithmetic. Instead one fixes a tolerance $\varepsilon > 0$ and defines the numerical rank as the number of singular values larger than ε. This matters in real option data, where bid–ask spreads, stale quotes, interpolation, and Monte Carlo noise destroy exact linear relations.

Accordingly, the most realistic use of the method is as a diagnostic:

- to detect redundant or nearly redundant instruments,
- to estimate the effective dimension of a static payoff family,
- and to compare model-implied payoff relations with observed market prices.

Viewed in light of the earlier sections, this is the finite-dimensional version of the same structural idea: one isolates the directions that persist and the directions that are effectively filtered out. Here that information is carried not by a Dirac kernel but by numerical rank, small singular values, and the kernel of the Gram matrix. If a portfolio lies approximately in $\ker T$ under the model but carries a non-negligible market price, then one has evidence of either mispricing, inconsistent quotes, or model misspecification. In that sense the Gram-matrix test is not merely an elementary linear-algebra exercise; it is the practical diagnostic shadow of the structural invariant viewpoint developed in the earlier sections.

10.5. K–theoretic interpretation

In a finite-dimensional setting, the operator-algebraic structure is much simpler than in the compactified torus model. The algebra generated by finitely many payoff vectors is finite-dimensional, and the orthogonal projection

$$P : \mathbb{R}^m \to \mathrm{span}\{x_0, \ldots, x_n\}$$

has rank

$$r = \mathrm{rank}(T).$$

This rank is locally stable under perturbations that do not push a singular value through zero. The relevant invariant therefore a stable integer measuring effective payoff-space dimension.

In this sense, Section 10 should be read as a finite-dimensional shadow of the earlier K–theoretic viewpoint. In the compactified geometric model, one studies kernel dimension, projection onto distinguished sectors, and stable index data. In the finite market, the corresponding objects are numerical rank, projection onto the payoff span, and kernel portfolios. The finite-dimensional calculation is therefore not a literal K–theory computation in the full operator-algebraic sense, but it captures the same structural question: how much independent payoff information is genuinely present, and which directions are redundant or nearly

11. Conclusions and outlook

We showed that the analytic machinery underlying classical diffusion-based option pricing—second–order elliptic operators, heat semigroups, and Gaussian kernels—runs in close parallel with the analytic machinery used in the heat–kernel approach to index theory. In simple one– and two–dimensional compactified models, we constructed Dirac–type operators whose squares reproduce the diffusion core of stylised pricing generators. In the twisted two–torus model, the associated Dirac operator has index k, and its zero modes can be written explicitly in terms of Jacobi theta functions.

Taken together, these results established three main structural points. First, stylised diffusion pricing generators on compact state spaces admit a natural Dirac–type realisation. Second, in the twisted two–torus model, the resulting operator has a topologically controlled index and an explicit theta–function zero–mode sector. Third, this zero–mode sector yields a finite-dimensional component of payoff space that remains fixed under the stylised backward diffusion semigroup, while the orthogonal complement is smoothed and damped by the heat flow.

From the finance side, the contribution should be read in a disciplined way. The present article did not propose a new empirical calibration framework, nor does it provide a complete implementable hedging methodology under discrete rebalancing, transaction costs, liquidity constraints, stochastic volatility, jumps, or other market frictions. Rather, it developed a structural operator-theoretic viewpoint on diffusion pricing generators and showed how this viewpoint led to invariant information about payoff spaces. The most accurate economic reading is therefore one of stylised dimensional reduction and spectral filtering: the geometric construction isolates a distinguished persistent sector, and the finite-market example of Section 10 shows that the same perspective has a practical finite-dimensional shadow through Gram-matrix/rank diagnostics for redundancy, constrained payoff-space dimension, and kernel portfolios.

Einstein's work on Brownian motion and the Stokes–Einstein relation (Einstein, 1905) provides a useful conceptual template for this picture. In physical diffusion, a single coefficient D summarises complicated microscopic randomness. In the present financial setting, the effective diffusion constant

$$D = \frac{\sigma^2}{2}$$

plays an analogous analytic role. The Dirac-type formulation then shows how this diffusion parameter interacts, in the compactified model, with the global geometry of state space and with the topology of the twisting bundle.

The spectral triple (A, H, D_A) provides a natural entry point for K–theory and noncommutative geometry into this discussion (Connes, 1994; Zois, 2000, 2010). In this framework, the zero-mode sector is the most visible analytic manifestation of an underlying index pairing. This does not yet produce a market rule by itself, but it does place the study of diffusion pricing generators into a setting where stable integer invariants can be defined and computed.

The main limitation of the present work was equally clear. The compactified circle and torus models are stylised geometric laboratories, not literal market models. Their purpose is to isolate the elliptic core of diffusion pricing in a setting where the spectral and topological structure becomes explicit. Thus the "stability" established here is mathematical stability of kernel dimension under smooth deformations within the stylised model, not automatic economic robustness in real markets. In particular, nothing in the article by itself guaranteed implementability under realistic trading constraints, nor does the zero–mode sector automatically survive unchanged once one leaves the idealised diffusion setting and introduces richer dynamics or market frictions. These features may well persist, at least in part, in more realistic models, but that is a matter for explicit analysis and practical verification rather than something established by the present article.

Several natural extensions remain open, and it is useful to distinguish between near-term continuations of the present article and longer-range research directions.

As immediate extensions , one may study higher-dimensional state spaces, more general covariance structures, and richer finite-market diagnostics tied to explicit instrument sets and calibrated examples. A particularly natural next step is to test, in calibrated or semi-calibrated models, how much of the zero-mode and dimensional-reduction picture survives once richer dynamics are introduced. It would also be valuable to supplement the present structural theory with more detailed feasibility and implementation analyses, including discrete rebalancing rules, numerical tolerances, and empirical or semi-empirical case studies.

At a more structural level, one may investigate odd-dimensional spectral refinements of the present framework, in particular *eta–invariant* type quantities measuring spectral asymmetry of nonzero modes in compact toy models (Atiyah et al., 1975a,b), as a complement to the index and zero-mode analysis developed here. More ambitiously, one may ask how far the present operator-theoretic viewpoint can be pushed beyond pure diffusion models to settings with stochastic volatility, jumps, other nonlocal effects or even noncommutative state spaces inspired by noncommutative tori.

Author contributions

The author is solely responsible for the conceptualization, methodology, formal analysis, investigation, writing—original draft preparation, and writing—review and editing of this manuscript.

Use of AI tools declaration

The author used AI tools for the matrix computations in subsection 10.3.

Acknowledgments

The author dedicates this article to the memory of his father, Panayiotis. The author wishes to thank Constantine Chronis from PPC for prompting his interest in this subject and for providing a convivial working ambience at PPC. The author also wishes to thank Sourin Ghosh from S&P Global and Andrew Lo from MIT for stimulating discussions and expositions related to this topic. No external funding was received for this research.

Conflict of interest

The author declares no conflicts of interest in this article.

Data availability

No empirical datasets were generated or analysed during the current study. The numerical toy example in Section 10 used only the hypothetical finite-market data reported in the article.

References

Alòs E, García-Lorite D (2025) On the Curvature of the Bachelier Implied Volatility. *Risks* 13: 27. https://doi.org/10.3390/risks13020027

Atiyah MF (1967) *K-Theory*. W. A. Benjamin, New York.

Atiyah MF, Singer IM (1963) The index of elliptic operators on compact manifolds. *Bull Amer Math Soc* 69: 422–433. https://doi.org/10.1090/S0002-9904-1963-10957-X

Atiyah MF, Patodi VK, Singer IM (1975a) Spectral asymmetry and Riemannian geometry I. *Math Proc Cambridge Philos Soc* 77: 43–69. https://doi.org/10.1017/S0305004100049410

Atiyah MF, Patodi VK, Singer IM (1975b) Spectral asymmetry and Riemannian geometry II. *Math Proc Cambridge Philos Soc* 78: 405–432. https://doi.org/10.1017/S0305004100051872

Bachelier L (1900) Théorie de la spéculation. *Ann Sci Ècole Norm Sup* 17: 21–86. https://doi.org/10.24033/asens.476

Berline N, Getzler E, Vergne M (1992) *Heat Kernels and Dirac Operators*. Springer, Berlin. https://doi.org/10.1007/978-3-642-58088-8

Bismut JM (1986) Index theorem and the heat equation. In: *Proceedings of the International Congress of Mathematicians* (Berkeley, 1986), 43–61.

Björk T (2009) *Arbitrage Theory in Continuous Time*. 3rd edition. Oxford University Press, Oxford.

Black F, Scholes M (1973) The pricing of options and corporate liabilities. *J Polit Econ* 81: 637–654. https://doi.org/10.1086/260062

Blackadar B (1998) *K-Theory for Operator Algebras*. 2nd edition. Cambridge University Press, Cambridge. https://doi.org/10.1017/9781009701907

Choi J, Kwak M, Tee CW, et al. (2022) A Black–Scholes user's guide to the Bachelier model. *J Futures Mark* 42: 959–980. https://doi.org/10.1002/fut.22315

Connes A (1994) *Noncommutative Geometry*. Academic Press, San Diego.

Einstein A (1905) On the movement of small particles suspended in a stationary liquid, as required by the molecular-kinetic theory of heat. *Ann Phys* 17: 549–560. https://doi.org/10.1002/andp.19053220806

Farkas HM, Kra I (1992) *Riemann Surfaces*. 2nd edition. Springer, New York. https://doi.org/10.1007/978-1-4612-2034-3

Freed DS (2021) The Atiyah–Singer index theorem. *Bull Amer Math Soc* 58: 495–532. https://doi.org/10.1090/bull/1747

Gatheral J (2006) *The Volatility Surface: A Practitioner's Guide*. Wiley, Hoboken. https://doi.org/10.1002/9781119202073

Gilkey PB (1984) *Invariance Theory, the Heat Equation, and the Atiyah–Singer Index Theorem*. Publish or Perish, Wilmington.

Griffiths P, Harris J (1978) *Principles of Algebraic Geometry*. Wiley, New York.

Halidias N (2025) An overview of financial mathematics with Python codes. *Monte Carlo Methods Appl* 31: 279–309. https://doi.org/10.1515/mcma-2025-2020

Higson N, Roe J (2000) *Analytic K-Homology*. Oxford Mathematical Monographs. Oxford University Press, Oxford. https://doi.org/10.1093/oso/9780198511762.001.0001

Karoubi M (1978) *K-Theory: An Introduction*. Springer-Verlag, Berlin. https://doi.org/10.1007/978-3-540-79890-5

Lawson HB Jr, Michelsohn ML (1989) *Spin Geometry*. Princeton University Press, Princeton.

Lindquist WB, Rachev ST, Gnawali J, et al. (2024) Dynamic asset pricing in a unified Bachelier–Black–Scholes–Merton model. *Risks* 12: 136. https://doi.org/10.3390/risks12090136

Merton RC (1973) Theory of rational option pricing. *Bell J Econ Manag Sci* 4: 141–183.

Reutter D (2015) The heat equation and the Atiyah–Singer index theorem. Essay, University of Oxford.

Roe J (1998) *Elliptic Operators, Topology and Asymptotic Methods*. 2nd edition. Longman, Harlow.

Schachermayer W, Teichmann J (2008) How Close Are the Option Pricing Formulas of Bachelier and Black–Merton–Scholes? *Math Financ* 18: 155–170. https://doi.org/10.1111/j.1467-9965.2007.00326.x

Shreve SE (2004) *Stochastic Calculus for Finance II: Continuous-Time Models*. Springer, New York.

Wegge-Olsen NE (1993) *K-Theory and C^*-Algebras: A Friendly Approach*. Oxford University Press, New York.

Zois IP (2000) A new invariant for σ models. *Commun Math Phys* 209: 757–783. https://doi.org/10.1007/PL00020964

Zois IP (2010) *18 Lectures on K-Theory*. arXiv:1008.1346 [math.KT].