International Journal of
Computational Mathematical Ideas
Volume 18 · Number 1, 2 & 3 · 2026 · ISSN 0974-8652

Available online at www.ijcmi.com

Modern AI Tools and the Indian IT Sector: A Jump-Augmented Physics-Informed Framework for Narrative Shocks, Repricing and Sector Allocation in NIFTY IT

S. Satyanarayana *

CEO & Chief Agentic AI Scientist, AlgoProfessor AI R&D Solutions, India

*Corresponding author: drssnaiml1@gmail.com

Article Info

Article history:
Received: 12-08-2026
Revised: 02-09-2026
Accepted: 08-09-2026
Available online: 2026

Keywords:
Modern AI tools; Agentic AI; Indian IT sector; NIFTY IT; Narrative shocks; Jump-diffusion; Physics-informed neural networks; Merton problem; Sector allocation; Backtesting.

Abstract

Between February and September 2026 the NIFTY IT index gave up roughly a quarter of its value while the underlying industry kept growing, a divergence that market commentary attributed almost entirely to the arrival of increasingly capable coding and agentic AI tools. This paper asks what the right allocation response to such a repricing is, and answers it with physics rather than pattern fitting. We extend the physics-informed pipeline of Part I from a pure diffusion to a jump-augmented Merton problem in which capability announcements arrive as a Poisson stream of two-sided, negatively skewed narrative shocks. The resulting Hamilton-Jacobi-Bellman equation is a partial integro-differential equation whose optimal fraction has no closed form; we prove that the first-order condition has a unique admissible root, and give an a-posteriori bound that controls the value-function error by the physics residual. A from-scratch physics-informed network, trained on the equation alone with no market data, recovers the semi-analytic value component to a relative L_2 error of 2.1×10^{-4} for a constant announcement intensity and 2.9×10^{-3} for a ramped one, with the realised error inside the bound by a factor of four. Calibrating to the observed NIFTY IT state of 10 September 2026 and backtesting over 400 simulated five-year paths net of 25 basis points of Indian transaction costs, we find a sharply asymmetric result. Averaged over states, jump awareness barely moves the static weight, from 0.569 to 0.546. Conditioned on the state it moves it by almost an order of magnitude, from 0.93 in quiet periods to 0.11 inside an announcement cluster. An oracle that is told the state reaches a Sharpe ratio of 0.39 against 0.28 for buy-and-hold, while a tradable trailing-window detector falls to 0.19 after estimation error and 4.3 units of annual turnover. The governing equation therefore supplies the correct response to AI narrative risk; the binding constraint is detecting the narrative state, and that is where engineering effort belongs. The backtest uses a calibrated market simulator, not real exchange data.

1. Introduction

Part I of this work brought physics-informed neural networks to portfolio optimisation and reached a deliberately unglamorous conclusion: the governing equation gives the correct allocation rule, and essentially all the residual value sits in estimating the rule's inputs [1]. That paper studied a pure diffusion. The Indian IT sector in 2026 has supplied a natural and pressing test of what happens when the diffusion assumption fails.

Over the first nine months of 2026 the NIFTY IT index fell by roughly a quarter, and the fall did not arrive as a diffusion. It arrived in clusters. In early February the index lost close to a fifth of its value in eight sessions, a speed of decline the sector had not seen since 2008, and the trigger was not an earnings miss but the release of agentic AI coding tools capable of automating work that Indian services firms have historically billed for [42]. Similar clusters followed each subsequent capability announcement [43, 44]. Meanwhile the industry itself kept growing: NASSCOM's Strategic

Review put FY26 technology-sector revenue at about \$315 billion, up 6.1 per cent, with headcount up only 2.3 per cent [39]. The market was not repricing realised cash flows. It was repricing a narrative about future ones.

That distinction matters for a control problem. A diffusion model says that risk accumulates smoothly and that the optimal risky weight depends on the premium-to-variance ratio alone. A market in which a discrete, unpredictable announcement can remove a fifth of sector value in eight sessions is not that market; it is a jump market, and the optimal policy in a jump market is different in kind, not merely in degree [12, 15]. The purpose of this paper is to write that difference down, solve it with the physics-informed machinery of Part I, and then measure honestly how much of the difference is capturable.

Our contributions are four.

1. We formulate sector allocation under *modern AI tool risk* as a jump-augmented Merton problem in which capability announcements arrive as a Poisson stream of two-sided, negatively skewed log-jumps, and we derive the reduced partial integro-differential HJB that the value function satisfies (Proposition 1).
2. We prove that the resulting first-order condition, which unlike the classical Merton condition admits no closed form, has a unique admissible root (Proposition 2), and we extend the a-posteriori residual bound of Part I to the jump case (Proposition 3), which licenses training on the physics residual alone.
3. We solve the equation with a from-scratch physics-informed network written in `numpy` with analytic gradients and no deep-learning library, and validate it against a semi-analytic reference to a relative L_2 error of 2.1×10^{-4} for a constant announcement intensity and 2.9×10^{-3} for a ramped one.
4. We calibrate to the observed NIFTY IT state of 10 September 2026 and report a cost-aware backtest whose lesson refines Part I's. Unconditionally, allowing for AI narrative jumps changes the static IT weight by only four per cent. Conditionally, it changes it by a factor of nearly nine. The whole of the value therefore lies in detecting the announcement state, and a naive trailing-window detector destroys that value after costs.

Throughout, the backtest runs on a calibrated simulator rather than exchange tick data, and we are explicit about that limitation in Section 9.

2. The 2026 repricing of Indian IT: stylised facts

Table 1 records the state of the NIFTY IT index on the reference date of this study, and Figure 1 shows the trailing-return profile that motivates the model. Three features stand out and each has a modelling consequence.

Table 1. Observed NIFTY IT state, 10 September 2026, as reported on the NSE index screen. Free-float market capitalisation is in lakh crore rupees. The derived rows are computed from the reported levels.

Quantity	Value	Quantity	Value
Index level	28,867.70	Trailing 1 week	-6.39%
Previous close	28,913.95	Trailing 1 month	-8.74%
Free-float market cap	10.77	Trailing 3 months	$+2.08\%$
Price / earnings	18.45	Trailing 6 months	-3.85%
Price / book	5.09	Year to date	-24.37%
52-week low	25,699.10	Trailing 1 year	-20.22%
52-week high	40,301.40	Trailing 3 years	-10.95%
Drawdown from 52-week high	-28.37%	*Recovery above 52-week low*	$+12.33\%$

The decline is discontinuous, not diffusive. A drawdown of -28 per cent from the 52-week high coexisting with a $+2$ per cent three-month return is not the profile of a steady grind. It is the profile of a small number of large adverse moves separated by ordinary trading. Analyst accounts of the period place those moves squarely on capability announcements from AI platform vendors rather than on results [40, 41]. A model of the sector must therefore carry a jump component whose arrivals are exogenous to the price process.

The shocks are two-sided but negatively skewed. Not every AI announcement is bad news for Indian IT. Capability releases that automate application development, testing and maintenance are adverse, and those segments are estimated to carry close to a third of industry revenue [40]. But AI-monetisation disclosures, deal wins and order-book expansion move the same names the other way, and several mid-tier firms outperformed through the correction [42]. This argues for a mixture jump law with a dominant adverse component and a smaller favourable one, rather than a one-sided disaster process.

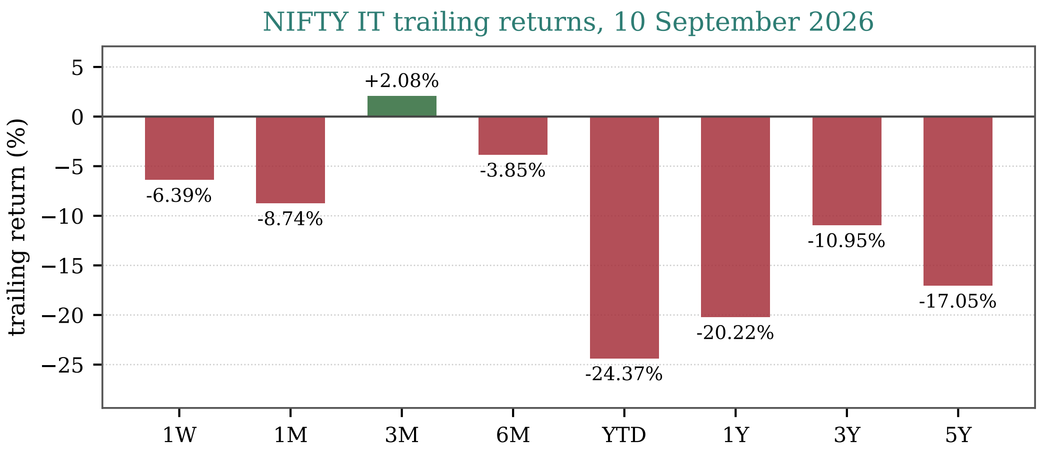

Figure 1. Trailing returns of the NIFTY IT index on the reference date. The term structure is the signature of a narrative repricing rather than a slow deterioration: the deep negatives sit at the year-to-date and one-year horizons that contain the February 2026 announcement cluster, while the three-month window, which contains a partial recovery, is positive.

Fundamentals and prices diverged. Sector revenue grew and AI-attributed revenue moved from pilot to line item, reaching an estimated \$10–12 billion industry-wide, while headcount growth stalled at 2.3 per cent [39]. Price-to-earnings compressed to 18.45, well below the sector's historical band. The repricing is thus a change in the market's assessment of the future revenue distribution, precisely the object that a jump intensity encodes. This is consistent with the broader evidence that capital markets reprice firms on measured exposure of their workforce to generative AI, and do so within days of a salient capability event [27, 29].

3. Background

3.1. Physics-informed learning

Physics-informed neural networks train a network to satisfy a differential equation by penalising the equation's residual at sampled collocation points together with initial and boundary conditions [2]. They sit inside a wider programme of physics-informed machine learning [3] supported by dedicated software [4] and by deep solvers for high-dimensional equations, including the deep Galerkin method [5], the deep BSDE method [6], and its formulation for nonlinear Black-Scholes and HJB problems [7]. The attraction for finance is the attraction in physics: the governing equation regularises the fit and reduces the appetite for data, which matters acutely when the regime of interest has only a handful of realisations.

3.2. Portfolio choice with jumps

Modern portfolio theory begins with the mean-variance trade-off [9]; its continuous-time counterpart is Merton's [10, 11], in which a constant relative risk aversion investor holds a fixed fraction of wealth in the risky asset. Merton himself introduced discontinuous returns to option pricing [12]; Kou's double exponential model gave the jump law analytic tractability [13], and affine jump-diffusions systematised the family [14]. The consequence for allocation was worked out by Liu, Longstaff and Pan, who showed that event risk makes the investor behave as if facing a short position in options and pushes the optimal weight materially below its diffusion value [15], and by Ait-Sahalia and co-authors for the multivariate case [16]. The mathematical foundations are standard [17], with financial applications collected in monograph form [18] and the jump-process machinery in [19].

3.3. Modern AI tools and equity repricing

The empirical literature on generative AI and firm value is young but consistent. Eisfeldt, Schubert and Zhang construct firm-level exposure from workforce task composition and document immediate repricing around the release of ChatGPT, with a long-short exposure portfolio earning economically large cumulative abnormal returns over the two following weeks [27, 28]. Eloundou and co-authors supply the task-exposure measurement that underlies it [29]. Event studies on corporate AI disclosures find significant abnormal returns with pronounced sectoral heterogeneity [34], and supply-chain spillovers have been documented as well [35]. Attention-based studies using search volume find a capital rotation away from incumbents facing displacement risk [36]. The macroeconomics is contested: Acemoglu argues the aggregate productivity gains are modest [33], while field evidence on assisted work reports substantial task-level gains [30, 31, 32]. For our purposes the contest is the point. When the fundamental effect is genuinely uncertain, price moves are driven by revisions to a shared narrative, arriving on announcements, which is a jump process by construction [38].

3.4. Volatility and the Indian market

Any credible simulator must respect the stylised facts of returns: fat tails, volatility clustering and leverage [21], captured by the ARCH and GARCH families [22, 23] and by regime-switching models [24]. Indian index data are well documented in this respect, with GARCH-class studies of NIFTY and SENSEX confirming strong persistence and asymmetry [25, 26]. Our simulator is a regime-switching jump diffusion calibrated to sit inside this documented range.

4. Problem formulation

An investor allocates a fraction π_t of wealth W_t to the IT index and the remainder to a risk-free rate r. The index is driven by a diffusion plus a compound Poisson stream of *narrative shocks* representing modern-AI-tool announcements,

$$\frac{\mathrm{d}S_t}{S_{t^-}} = \mu\,\mathrm{d}t + \sigma\,\mathrm{d}B_t + \mathrm{d}\left(\sum_{k=1}^{N_t} y_k\right), \qquad y_k = e^{Z_k} - 1, \tag{1}$$

where N_t is a Poisson process with deterministic intensity $\lambda(t)$ and the log-jumps Z_k are independent draws from the two-component mixture

$$Z \sim w_\downarrow \mathcal{N}\left(m_\downarrow, \delta_\downarrow^2\right) + \left(1 - w_\downarrow\right) \mathcal{N}\left(m_\uparrow, \delta_\uparrow^2\right), \qquad m_\downarrow < 0 < m_\uparrow. \tag{2}$$

The adverse component represents capability releases that automate billable services work; the favourable component represents AI-monetisation and deal-win news. Writing $\lambda(t)$ as time-varying is what lets the model express an announcement cycle: intensity is low in quiet stretches and high while a capability cycle is in progress.

Wealth evolves as $\mathrm{d}W_t = W_{t^-}[(r + \pi_t(\mu - r))\,\mathrm{d}t + \pi_t\sigma\,\mathrm{d}B_t + \pi_t\,\mathrm{d}(\sum_k y_k)]$, and the investor maximises $\mathbb{E}[U(W_T)]$ under constant relative risk aversion $U(w) = w^{1-\gamma}/(1-\gamma)$, $\gamma > 0$, $\gamma \neq 1$. We restrict $\pi \in [0,1]$: no leverage and no short sales, which is both the realistic institutional constraint and sufficient to keep $1 + \pi y > 0$ almost surely.

Proposition 1 (Reduced integro-differential HJB)**.** *The value function separates as $V(t,w) = \varphi(t)U(w)$, where φ solves the linear terminal-value problem*

$$\varphi'(t) + \rho(t)\varphi(t) = 0, \qquad \varphi(T) = 1, \tag{3}$$

with

$$\rho(t) = (1-\gamma)\left[r + \pi^*(t)(\mu - r) - \tfrac{1}{2}\gamma\pi^*(t)^2\sigma^2\right] + \lambda(t)\mathbb{E}\left[\left(1 + \pi^*(t)y\right)^{1-\gamma} - 1\right], \tag{4}$$

so that $\varphi(t) = \exp\left(\int_t^T \rho(u)\,\mathrm{d}u\right)$, and the optimal fraction $\pi^*(t)$ is the root in $[0,1]$ of the first-order condition

$$g(\pi;\lambda(t)) := (\mu - r) - \gamma\pi\sigma^2 + \lambda(t)\mathbb{E}\left[y(1+\pi y)^{-\gamma}\right] = 0. \tag{5}$$

Proof. Substituting the ansatz $V(t,w) = \varphi(t)w^{1-\gamma}/(1-\gamma)$ into the HJB equation

$$V_t + \max_\pi \left\{ (r + \pi(\mu - r))\,wV_w + \tfrac{1}{2}\pi^2\sigma^2 w^2 V_{ww} + \lambda(t)\mathbb{E}\left[V(t, w(1+\pi y)) - V(t,w)\right] \right\} = 0$$

and using $wV_w = (1-\gamma)V$, $w^2V_{ww} = -\gamma(1-\gamma)V$ and $\mathbb{E}[V(t,w(1+\pi y)) - V(t,w)] = V(t,w)\,\mathbb{E}[(1+\pi y)^{1-\gamma} - 1]$, every term carries the common factor $w^{1-\gamma}/(1-\gamma)$. Dividing it out gives $\varphi' + \rho\varphi = 0$ with ρ as in (4) once the maximising π is substituted. Differentiating the inner objective in π and cancelling the factor $(1-\gamma)\varphi$, which is nonzero, yields (5). The terminal condition $V(T,w) = U(w)$ gives $\varphi(T) = 1$, and integrating the linear ordinary differential equation backwards from T gives the exponential form. □

Equation (5) is where the extension bites. In Part I the corresponding condition was linear in π and gave the closed form $\pi^* = (\mu - r)/(\gamma\sigma^2)$. Here the jump expectation makes it transcendental, so the policy must be computed numerically. Proposition 2 shows the computation is well posed.

Proposition 2 (Existence and uniqueness of the admissible root)**.** *Fix $\lambda \geq 0$, $\gamma > 0$ and $\sigma > 0$, and suppose $\mathbb{E}|y| < \infty$ with $\mathbb{P}(y < 0) > 0$. Then on $\pi \in [0,1]$ the map $\pi \mapsto g(\pi;\lambda)$ is continuous and strictly decreasing, so (5) has at most one interior root, and exactly one if $g(0;\lambda) > 0 > g(1;\lambda)$; otherwise the maximiser is the corresponding endpoint. The resulting π^* is non-increasing in λ whenever $\mathbb{E}[y(1+\pi y)^{-\gamma}] < 0$.*

Proof. Differentiating under the expectation, which is justified by dominated convergence on the compact set $\pi \in [0,1]$ with $1+\pi y$ bounded away from zero,

$$g'(\pi;\lambda) = -\gamma\sigma^2 - \lambda\gamma\mathbb{E}\left[y^2(1+\pi y)^{-\gamma-1}\right] < 0,$$

since $\gamma\sigma^2 > 0$ and the expectation is non-negative. Strict monotonicity gives uniqueness, and continuity with a sign change gives existence by the intermediate value theorem; if no sign change occurs, g has a constant sign and the concave objective is maximised at an endpoint. For the comparative static, $\partial g/\partial\lambda = \mathbb{E}[y(1+\pi y)^{-\gamma}]$, so by the implicit function theorem $\partial\pi^*/\partial\lambda = -(\partial g/\partial\lambda)/g'(\pi^*)$, which is non-positive exactly when $\mathbb{E}[y(1+\pi y)^{-\gamma}] \le 0$. □

Remark 1. The final clause of Proposition 2 is the economic content of the model. The sign of $\mathbb{E}[y(1+\pi y)^{-\gamma}]$ is the risk-adjusted mean shock: the utility weighting $(1+\pi y)^{-\gamma}$ penalises adverse jumps more heavily than it rewards favourable ones of equal magnitude, so even a jump law with mean close to zero produces a negative risk-adjusted mean and hence a policy that de-risks as announcements become more frequent. Modern AI tools do not need to be net-destructive of Indian IT earnings for the optimal weight to fall; it is enough that they make the outcome distribution more discontinuous and left-skewed.

Proposition 3 (A-posteriori residual bound, jump case)**.** *Let φ solve (3) with ρ given by (4), and let $\hat{\varphi}$ be any C^1 approximation with residual $R(t) := \hat{\varphi}'(t) + \rho(t)\hat{\varphi}(t)$ and terminal error $\eta := \hat{\varphi}(T) - 1$. Then for all $t \in [0,T]$,*

$$|\hat{\varphi}(t) - \varphi(t)| \;\le\; \left(|\eta| + \int_t^T |R(u)|\,\mathrm{d}u\right)\exp\left(\int_t^T |\rho(u)|\,\mathrm{d}u\right). \tag{6}$$

Proof. The error $e = \hat{\varphi} - \varphi$ satisfies $e'(t) + \rho(t)e(t) = R(t)$ with $e(T) = \eta$. Integrating backwards from T and applying the triangle inequality gives $|e(t)| \le |\eta| + \int_t^T |R| + \int_t^T |\rho||e|$, and Gronwall's inequality yields (6). The jump term enters only through ρ, which is bounded on $[0,T]$ whenever λ is bounded and $\mathbb{E}[(1+\pi y)^{1-\gamma}]$ is finite, so the bound holds verbatim for the integro-differential problem. □

Proposition 3 is the licence for physics-informed training: the value-function error is controlled by the residual we penalise plus the terminal error we penalise, so driving both to zero drives the solution to the truth. Proposition 2 supplies the second guarantee, namely that the policy extracted from the first-order condition is the unique admissible one.

5. Method

5.1. The network and its loss

We approximate $\varphi_\theta(\tau)$ on the scaled horizon $\tau = t/T \in [0,1]$ with a small single-hidden-layer tanh network, $\varphi_\theta(\tau) = a_0 + \sum_j v_j \tanh(w_j\tau + b_j)$, whose time derivative $\varphi'_\theta(\tau) = \sum_j v_j w_j \operatorname{sech}^2(w_j\tau + b_j)$ is available in closed form. Training minimises the physics loss

$$\mathscr{L}(\theta) = \frac{1}{N}\sum_{i=1}^{N} R(\tau_i)^2 + \Lambda\left(\varphi_\theta(1) - 1\right)^2, \qquad R = \varphi'_\theta + T\rho\,\varphi_\theta, \tag{7}$$

over N collocation points with Adam [8]. The parameter gradients of the residual are derived analytically, so no automatic differentiation and no deep-learning library is used; the implementation is `numpy` only. The jump expectations in ρ and in the first-order condition are evaluated by Gauss-Hermite quadrature on each mixture component, which is exact for the Gaussian log-jump law up to the node count and avoids Monte-Carlo noise inside the training loop. Figure 2 shows the structure.

5.2. The AI-narrative pipeline

Figure 3 places the solver inside an end-to-end pipeline that differs from Part I's in one important respect: the calibration stage now has two channels. The conventional channel estimates drift and volatility from index history. The second channel estimates the announcement process itself, from the arrival record of capability releases and from the market's own jump statistics, and it is this channel that the backtest of Section 8 shows to be decisive.

5.3. Procedure

Figure 4 states the procedure in full.

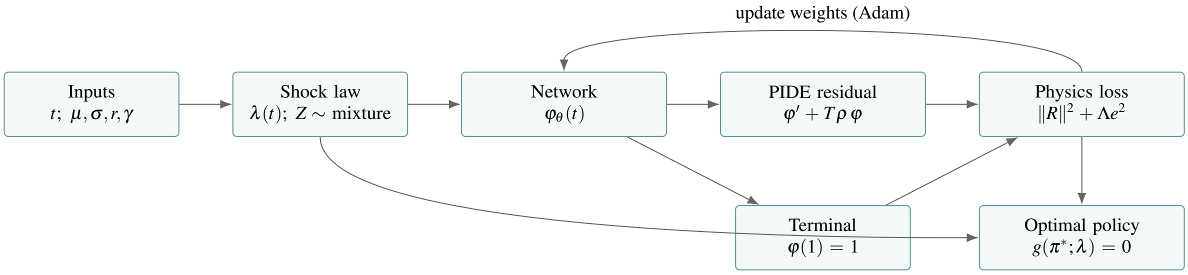

Figure 2. The jump-augmented physics-informed solver. Market parameters and the announcement-shock law feed the network φ_θ; its output and analytic derivative form the integro-differential residual and the terminal condition, whose weighted sum is the physics loss; the allocation follows from the first-order condition of Proposition 1, whose unique admissible root is guaranteed by Proposition 2.

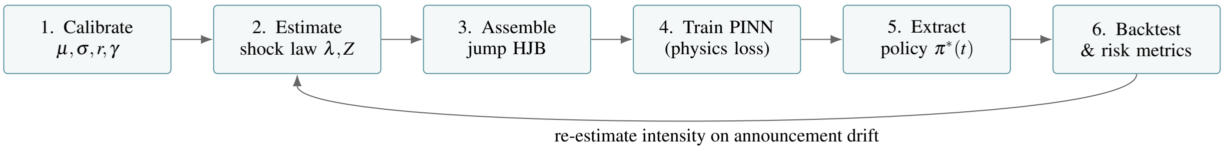

Figure 3. The AI-narrative allocation pipeline. Relative to Part I the calibration stage splits into a conventional moment channel and an announcement-process channel, and the feedback loop closes on the intensity estimate rather than on the drift.

Algorithm 1: jump-aware physics-informed sector allocation

Input: calibration μ, σ, r, γ; shock law $\lambda(\cdot)$, $(w_\downarrow, m_\downarrow, \delta_\downarrow, m_\uparrow, \delta_\uparrow)$; horizon T; collocation budget N.

1. build Gauss-Hermite nodes and weights for each mixture component of Z;
2. for each τ on the collocation grid, solve $g(\pi; \lambda(\tau)) = 0$ by safeguarded Newton bracketed to $[0, 1]$ (Proposition 2) to obtain $\pi^*(\tau)$;
3. assemble the reduced rate $\rho(\tau)$ from (4);
4. train φ_θ to minimise (7) with Adam and analytic gradients;
5. verify the solve against the semi-analytic reference and against the a-posteriori bound (6);
6. backtest π^* on simulated index paths net of costs, alongside buy-and-hold, a fixed mix, a jump-blind Merton rule, an oracle and a causal detector;
7. report CAGR, volatility, Sharpe, Sortino, maximum drawdown, tail loss and turnover;
8. adaptation loop: monitor realised against assumed announcement intensity; on drift, re-estimate the shock law and return to step 2.

Output: allocation policy $\pi^*(\cdot)$ and risk report.

Figure 4. The jump-aware allocation and backtest loop. Step 2 is the policy solve guaranteed by Proposition 2; step 4 is the PINN solve of Proposition 1, whose accuracy is controlled by Proposition 3; steps 6 to 8 evaluate and adapt it.

6. Calibration

Table 2 lists the calibration. The risk-free rate sits in the Indian policy-rate neighbourhood. The diffusive drift and volatility are set so that, once the jump component is added, the index has an unconditional expected return of 14.16 per cent and an unconditional volatility of 21.86 per cent, a NIFTY-IT-like pair. Risk aversion is set at $\gamma = 3$, a standard value for a single-sector allocation.

The shock law is the part that carries the paper's subject matter. Announcements arrive at 1.2 per year in quiet stretches and 5.0 per year inside an announcement cycle; with the monthly persistence probabilities in Table 2, the stationary probability of the cycle state is one third and the average intensity is $\bar{\lambda} = 2.47$ per year. Sixty-five per cent of announcements are adverse with a mean log-impact of -7.5 per cent and dispersion 7.0 per cent, and the remainder are favourable with mean $+6.0$ per cent and dispersion 3.0 per cent. This produces a mean shock of $\mathbb{E}[y] = -2.37$ per cent and a second moment of $\mathbb{E}[y^2] = 7.66 \times 10^{-3}$, so jumps account for 39.5 per cent of total index variance while contributing only a modest drag to the mean. Round-trip cost is 25 basis points per unit of weight change, consistent with Indian securities transaction tax, brokerage and slippage.

Table 2. Calibration and the physics-informed solve. The jump component supplies about two fifths of index variance while shifting the mean only slightly, which is the configuration under which the state-dependence result of Section 7 is sharpest.

Quantity	Value
Risk-free rate r	6.0%
Diffusive drift μ	20.0%
Diffusive volatility σ	17.0%
Risk aversion γ	3.0
Horizon T	5 years
Announcement intensity, quiet state λ_q	1.2 / year
Announcement intensity, cycle state λ_s	5.0 / year
Stationary probability of cycle state	0.333
Average intensity $\bar{\lambda}$	2.47 / year
Adverse share $w_\downarrow$	0.65
Adverse log-jump $(m_\downarrow, \delta_\downarrow)$	$(-7.5\%, 7.0\%)$
Favourable log-jump $(m_\uparrow, \delta_\uparrow)$	$(+6.0\%, 3.0\%)$
Implied mean shock $\mathbb{E}[y]$	-2.37%
Implied second moment $\mathbb{E}[y^2]$	7.66×10^{-3}
Unconditional return $\mu + \bar{\lambda}\mathbb{E}[y]$	14.16%
Unconditional volatility	21.86%
Jump share of total variance	39.5%
Jump-blind Merton fraction	0.5691
Jump-aware fraction $\pi^*(\bar{\lambda})$	0.5459
Jump-aware fraction, quiet state	0.9288
Jump-aware fraction, cycle state	0.1061
PINN error on φ, constant λ (rel. L_2)	2.1×10^{-4}
PINN error on φ, ramped $\lambda(t)$ (rel. L_2)	2.9×10^{-3}
Round-trip transaction cost	25 bps

7. The network solves the physics

Figure 5 shows the recovered value component $\varphi(t)$ against the semi-analytic reference $\exp(\int_t^T \rho)$, evaluated by quadrature on the same grid. With no market data at all, only the equation, the network matches the reference to a relative L_2 error of 2.1×10^{-4} for a constant announcement intensity; at the initial time it predicts $\varphi(0) = 0.437778$ against the reference 0.437776. For the harder case, in which the intensity ramps from the quiet level into a cycle and then partially decays, the error rises to 2.9×10^{-3}, still a faithful solve of a problem whose rate function has two sharp transitions [5].

The a-posteriori bound of Proposition 3 is checked directly on the ramped case. The realised terminal error is $\eta = 2.05 \times 10^{-3}$ and the integrated residual gives a bound of 9.36×10^{-3}, against a realised maximum error of 2.43×10^{-3}. The bound holds and is tight to within a factor of four, which is the practically useful regime: it is informative enough to serve as a stopping criterion and conservative enough to be trusted. This is what makes residual-only training defensible. We never see the true solution during training, and we do not need to.

8. Results

8.1. The policy responds to shock frequency, not to shock average

Figure 6 and the heat map of Figure 7 display the comparative statics of Proposition 2. Two readings follow, and they point in opposite directions.

Averaged over states, jump awareness barely matters. The jump-blind Merton rule, formed from the same observable unconditional moments, sets the IT weight at 0.5691; the jump-aware rule at the average intensity sets it at 0.5459, a reduction of 4.2 per cent. An investor who estimated drift and volatility from index history and applied the classical formula would be very nearly right on average. This is a negative result and we report it as one.

Conditioned on the announcement state, jump awareness matters enormously. In a quiet stretch, where announcements arrive 1.2 times a year, the optimal weight is 0.9288. Inside an announcement cycle, where they arrive 5.0 times a year, it is 0.1061. The optimal policy moves by a factor of nearly nine across states that the unconditional average smooths away entirely. The mechanism is the one identified in the remark after Proposition 2: it is not that the expected shock

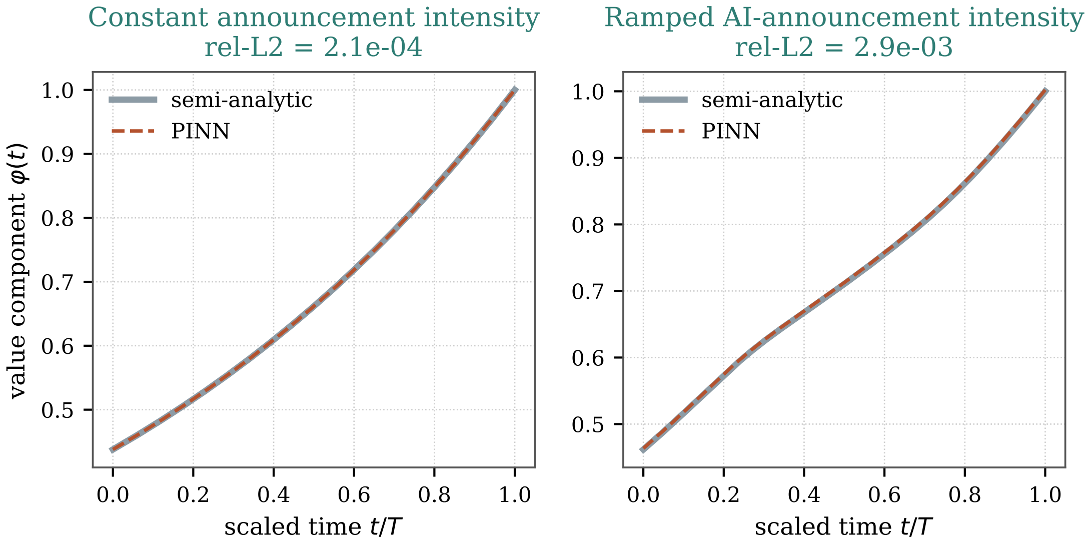

Figure 5. The physics-informed network (dashed) recovers the semi-analytic value component $\varphi(t)$ (solid) with no market data. Left: a constant announcement intensity, relative L_2 error 2.1×10^{-4}. Right: an intensity that ramps into an announcement cycle and partially decays, relative L_2 error 2.9×10^{-3}.

becomes much worse in a cycle, since the jump law itself is unchanged, but that the risk-adjusted shock cost scales linearly with arrival frequency, and the utility weighting $(1+\pi y)^{-\gamma}$ makes that cost bite hardest exactly where the position is largest.

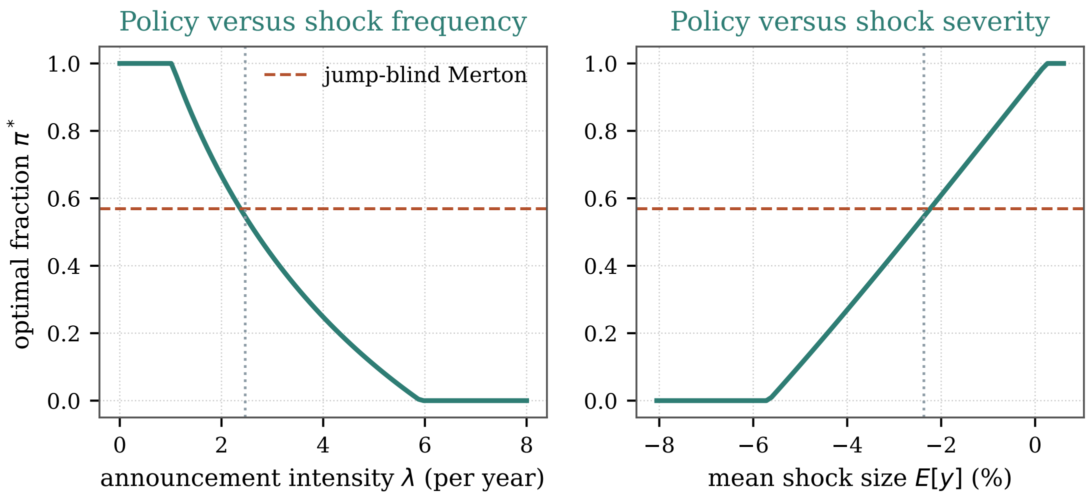

Figure 6. Optimal IT weight against announcement frequency (left) and against the mean shock size (right), with the jump-blind Merton weight shown for reference and the calibration point marked. The policy is far more sensitive to how often announcements arrive than to how bad the average one is.

8.2. Backtest

Table 3 and Figures 8 to 10 report the backtest over 400 simulated five-year paths, net of 25 basis points of cost. Six strategies are compared: buy-and-hold on the IT index, a fixed sixty-forty mix, the jump-blind Merton rule, the jump-aware constant rule that the PINN pipeline produces, an unrealisable oracle that is told the announcement state at each step, and a tradable causal detector that estimates intensity from a trailing sixty-three-day window by counting returns beyond 2.5 rolling standard deviations.

Buy-and-hold earns the highest raw return at 12.68 per cent per year, but carries 22.16 per cent volatility and a −33.91 per cent worst drawdown, so its Sharpe ratio is the lowest of the static strategies at 0.276 [20]. The three static de-risked policies cluster tightly: the fixed mix at 0.326, the jump-blind Merton rule at 0.331, and the jump-aware rule at 0.334,

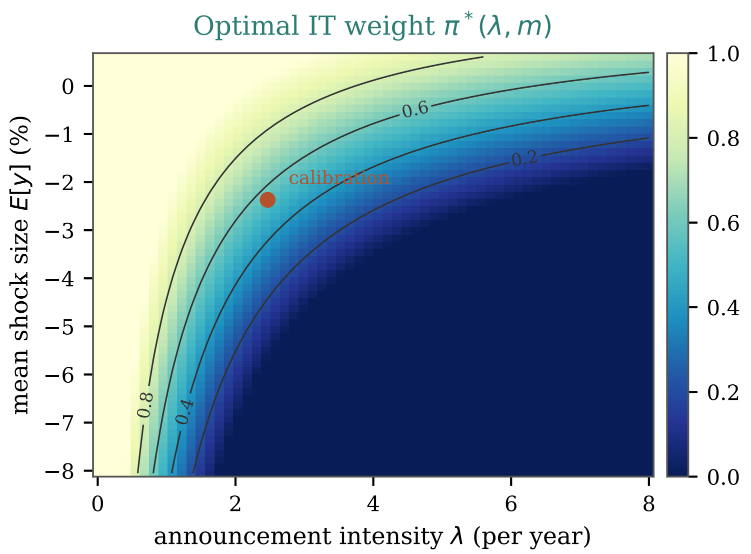

Figure 7. The optimal IT weight over the announcement-intensity and mean-shock plane. Contours are near-vertical over most of the region, which is the geometric statement of the same finding: frequency dominates severity.

with the jump-aware rule giving the shallowest static drawdown at −18.16 per cent and the smallest daily tail loss. The improvement from jump awareness is real but small, which is exactly what the 4.2 per cent weight difference predicts.

The oracle is the ceiling and it is a high one: Sharpe 0.387, achieved by holding 0.93 in quiet periods and 0.11 inside announcement cycles at a cost of 1.25 units of annual turnover. The gap between 0.387 and 0.334 is the value of solving the announcement-detection problem. The causal detector attempts to capture it and does the opposite: after estimation error and 4.26 units of annual turnover at 25 basis points, its Sharpe falls to 0.194, below every static rule including buy-and-hold, with a −25.55 per cent drawdown. A detector that reacts to realised volatility is structurally late, because in a jump market the information arrives in the jump itself. It sells after the announcement has already repriced the index and buys back after the recovery has already happened.

Table 3. Backtest over 400 simulated paths, five years, net of 25 bps costs. Sharpe and Sortino are computed on log growth so that they distinguish constant-weight policies, which a simple-return Sharpe ratio cannot. CVaR is the mean daily loss in the worst five per cent of sessions.

Strategy	CAGR	Vol	Sharpe	Max DD	CVaR_{95}	Turn/yr
Buy-and-hold (NIFTY IT)	12.68%	22.16%	0.276	−33.91%	−2.87%	0.00
Fixed 60/40	10.56%	13.19%	0.326	−20.15%	−1.71%	0.00
Merton, jump-blind	10.37%	12.50%	0.331	−19.01%	−1.62%	0.00
PINN/Merton, jump-aware	10.22%	11.99%	0.334	−18.16%	−1.56%	0.00
Jump-aware (causal est.)	9.69%	16.93%	0.194	−25.55%	−2.32%	4.26
Jump-aware (oracle)	12.47%	15.06%	**0.387**	−22.04%	−2.05%	1.25

9. Discussion

The experiments support a reading that is more specific than Part I's and, we believe, more useful. Part I concluded that the physics gives the rule and the value is in the inputs. This paper locates which input. It is not the equity premium, whose estimation error was Part I's villain, and it is not the diffusive volatility, which a practitioner estimates adequately from history. It is the arrival intensity of narrative shocks.

That conclusion has an uncomfortable corollary for the 2026 episode. A sector that has just absorbed a −24 per cent year-to-date move on announcement risk is, under this model, either cheap or correctly priced depending entirely on a quantity nobody observes directly, namely how frequently the next capability release will land. The price-to-earnings ratio of 18.45 in Table 1 is a statement about the market's implied intensity, not about earnings. The forecasting literature is instructive here: markets reprice AI exposure within days of a salient event [27], but the underlying productivity question is genuinely unsettled, with credible aggregate estimates far below the task-level ones [33, 30]. A model that admits it cannot forecast the intensity is being honest about the state of knowledge, not evading the question.

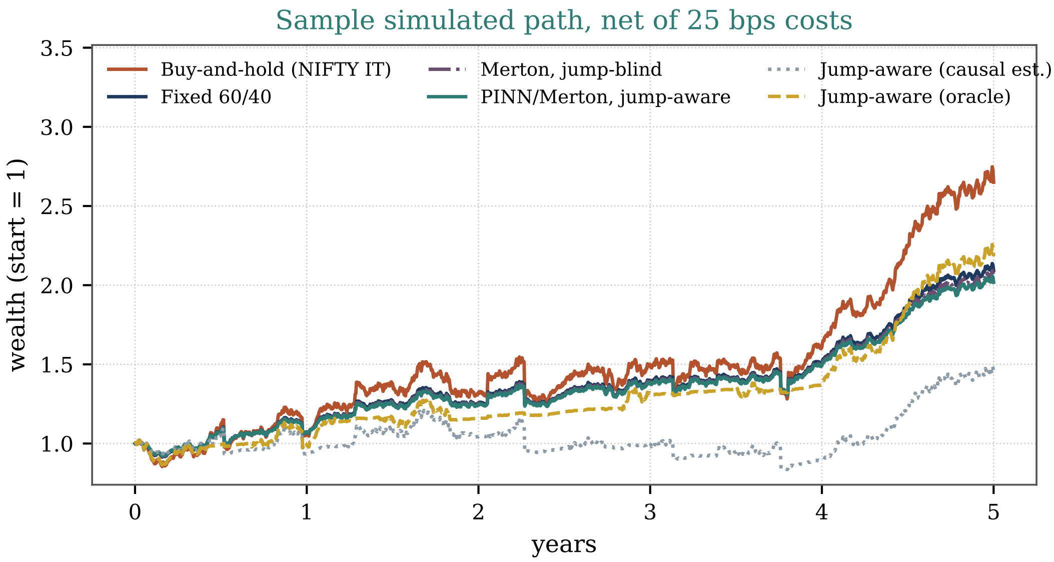

Figure 8. A representative simulated wealth path, net of costs. The oracle tracks buy-and-hold on the way up and steps aside through announcement clusters; the static de-risked policies give up terminal wealth for a much smoother ride; the causal detector lags after costs.

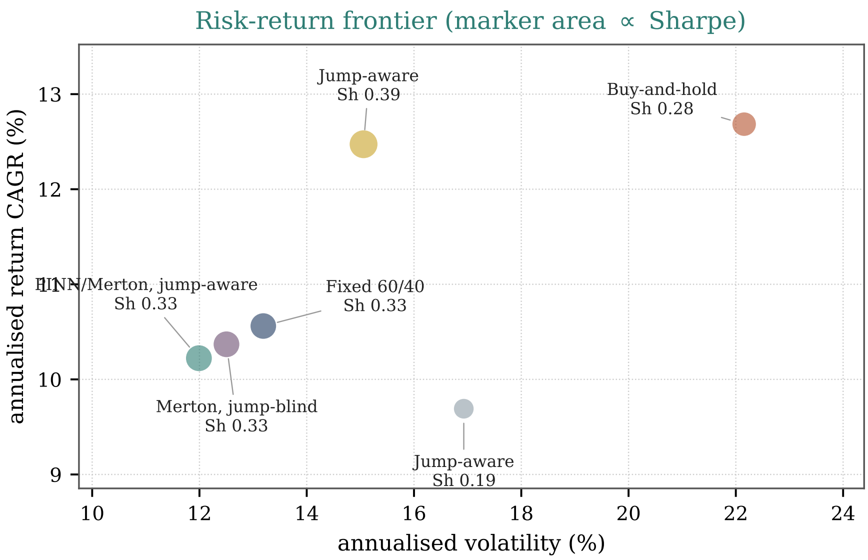

Figure 9. Risk-return positions averaged over paths, with marker area proportional to Sharpe. The jump-aware static policy sits at the best risk-adjusted point among the tradable static rules; the oracle is the ceiling; the causal detector is dominated.

Three practical implications follow. First, effort in a production system should move from premium estimation to announcement-process estimation: release calendars, vendor roadmaps, patent and hiring signals, and options-implied jump risk are all more informative about λ than a trailing return window is. Second, any adaptive rule must be tested against its turnover bill before it is believed; the causal detector here loses 0.14 of Sharpe to estimation error and roughly the same again to costs. Third, the near-vertical contours of Figure 7 mean that a coarse but timely intensity signal is worth more than a precise but lagged severity estimate, which is an unusual and actionable ordering of research priorities.

Limitations. The backtest is a calibrated simulation, not real NSE or BSE data, so the absolute figures should be read as internal comparisons rather than forecasts [25]. The model has a single risky asset, whereas the observed correction was highly heterogeneous across large-cap and mid-tier names, and the cross-sectional dispersion that a multi-asset formulation would capture is exactly where the reported outperformance of some mid-tier firms lives [42]. Volatility is constant between jumps, whereas Indian index returns display well-documented clustering and asymmetry that a

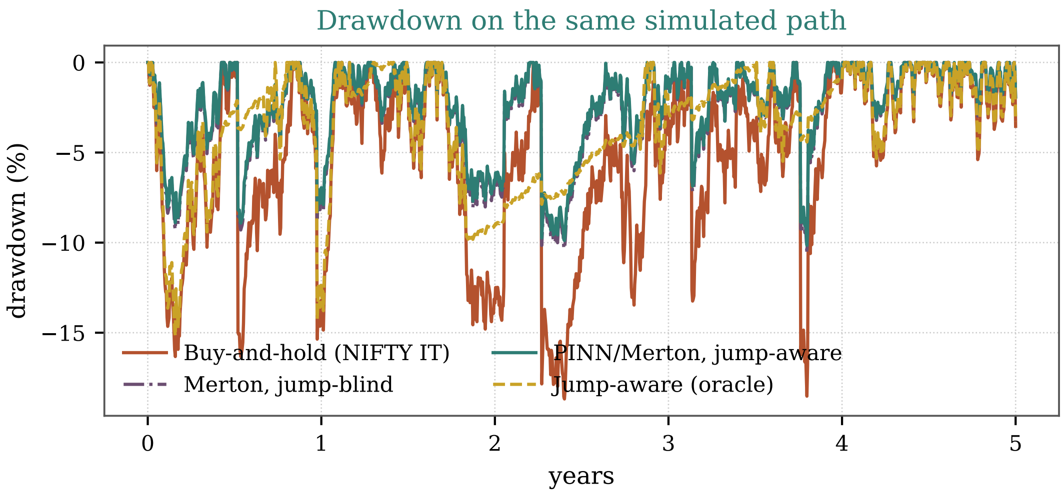

Figure 10. Drawdown on the same simulated path. Announcement clusters appear as near-vertical steps rather than gradual descents, which is precisely why a volatility-triggered detector arrives too late to avoid them.

GARCH layer would capture [23, 26]. The announcement intensity is exogenous and deterministic; a self-exciting Hawkes specification, in which announcements beget announcements, would fit the observed clustering better and is the natural next step [19]. Finally, constant relative risk aversion is what makes the value function separate and the problem one-dimensional; richer utilities, intermediate consumption, or transaction-cost-aware control break that separation and give a genuinely multidimensional integro-differential HJB, which is where deep solvers earn their keep [6].

Future work runs in three directions: a Hawkes-driven intensity estimated jointly from price jumps and an announcement calendar; a multi-asset formulation over the NIFTY IT constituents with exposure-weighted jump loadings in the spirit of firm-level AI exposure measurement [27, 29]; and treating intensity estimation as a first-class regularised learning problem so that adaptation helps rather than hurts [37].

10. Conclusion

We extended the physics-informed allocation pipeline of Part I from a diffusion to a jump market in order to address a concrete question: what is the correct portfolio response to the repricing of Indian IT by modern AI tools. The extension required a genuine mathematical step, since the optimal fraction under jumps has no closed form, and we supplied the two guarantees that make the numerical solve trustworthy: uniqueness of the admissible root of the first-order condition, and an a-posteriori bound that controls value-function error by the physics residual. A from-scratch network trained on the equation alone recovered the value component to 2.1×10^{-4} for a constant announcement intensity and 2.9×10^{-3} for a ramped one, with realised error inside the bound.

The empirical lesson is sharply asymmetric and worth stating plainly. Unconditionally, jump awareness is nearly irrelevant: it moves the static IT weight from 0.569 to 0.546 and lifts Sharpe from 0.331 to 0.334. Conditionally, it is decisive: the optimal weight swings from 0.93 in quiet stretches to 0.11 inside an announcement cluster, and an oracle that knows which state it is in earns 0.387 against 0.276 for buy-and-hold. A tradable trailing-window detector captures none of that and finishes at 0.194 after costs, because in a jump market the information arrives inside the jump. Embedding the governing equation therefore tells us exactly what to do about AI narrative risk once we know how often it strikes; it does not tell us how often it will strike, and honest system design should place its effort there.

Acknowledgements

The author thanks colleagues at AlgoProfessor AI R&D Solutions for helpful discussion. All code needed to reproduce every figure, table and number in this paper is available from the author. Index levels and trailing returns in Table 1 and Figure 1 are as reported on the NSE index screen for 10 September 2026; the backtest uses a seeded market simulator calibrated to those observations and does not use real exchange data.

References

[1] S. Satyanarayana, "Physics-informed neural networks for portfolio optimisation and algorithmic trading: a reproducible pipeline with an Indian-market backtest," *International Journal of Computational Mathematical Ideas*, vol. 18, no. 1, pp.

202600180088–202600180095, 2026.

[2] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," *Journal of Computational Physics*, vol. 378, pp. 686–707, 2019.

[3] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang, "Physics-informed machine learning," *Nature Reviews Physics*, vol. 3, no. 6, pp. 422–440, 2021.

[4] L. Lu, X. Meng, Z. Mao, and G. E. Karniadakis, "DeepXDE: a deep learning library for solving differential equations," *SIAM Review*, vol. 63, no. 1, pp. 208–228, 2021.

[5] J. Sirignano and K. Spiliopoulos, "DGM: a deep learning algorithm for solving partial differential equations," *Journal of Computational Physics*, vol. 375, pp. 1339–1364, 2018.

[6] J. Han, A. Jentzen, and W. E, "Solving high-dimensional partial differential equations using deep learning," *Proceedings of the National Academy of Sciences*, vol. 115, no. 34, pp. 8505–8510, 2018.

[7] W. E, J. Han, and A. Jentzen, "Deep learning-based numerical methods for high-dimensional parabolic partial differential equations and backward stochastic differential equations," *Communications in Mathematics and Statistics*, vol. 5, no. 4, pp. 349–380, 2017.

[8] D. P. Kingma and J. Ba, "Adam: a method for stochastic optimization," in *Proc. International Conference on Learning Representations (ICLR)*, 2015.

[9] H. Markowitz, "Portfolio selection," *The Journal of Finance*, vol. 7, no. 1, pp. 77–91, 1952.

[10] R. C. Merton, "Lifetime portfolio selection under uncertainty: the continuous-time case," *The Review of Economics and Statistics*, vol. 51, no. 3, pp. 247–257, 1969.

[11] R. C. Merton, "Optimum consumption and portfolio rules in a continuous-time model," *Journal of Economic Theory*, vol. 3, no. 4, pp. 373–413, 1971.

[12] R. C. Merton, "Option pricing when underlying stock returns are discontinuous," *Journal of Financial Economics*, vol. 3, no. 1–2, pp. 125–144, 1976.

[13] S. G. Kou, "A jump-diffusion model for option pricing," *Management Science*, vol. 48, no. 8, pp. 1086–1101, 2002.

[14] D. Duffie, J. Pan, and K. Singleton, "Transform analysis and asset pricing for affine jump-diffusions," *Econometrica*, vol. 68, no. 6, pp. 1343–1376, 2000.

[15] J. Liu, F. A. Longstaff, and J. Pan, "Dynamic asset allocation with event risk," *The Journal of Finance*, vol. 58, no. 1, pp. 231–259, 2003.

[16] Y. Aït-Sahalia, J. Cacho-Diaz, and T. R. Hurd, "Portfolio choice with jumps: a closed-form solution," *The Annals of Applied Probability*, vol. 19, no. 2, pp. 556–584, 2009.

[17] W. H. Fleming and H. M. Soner, *Controlled Markov Processes and Viscosity Solutions*, 2nd ed. New York, NY: Springer, 2006.

[18] H. Pham, *Continuous-time Stochastic Control and Optimization with Financial Applications*. Berlin, Germany: Springer, 2009.

[19] R. Cont and P. Tankov, *Financial Modelling with Jump Processes*. Boca Raton, FL: Chapman & Hall/CRC, 2004.

[20] W. F. Sharpe, "The Sharpe ratio," *The Journal of Portfolio Management*, vol. 21, no. 1, pp. 49–58, 1994.

[21] R. Cont, "Empirical properties of asset returns: stylized facts and statistical issues," *Quantitative Finance*, vol. 1, no. 2, pp. 223–236, 2001.

[22] R. F. Engle, "Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation," *Econometrica*, vol. 50, no. 4, pp. 987–1007, 1982.

[23] T. Bollerslev, "Generalized autoregressive conditional heteroskedasticity," *Journal of Econometrics*, vol. 31, no. 3, pp. 307–327, 1986.

[24] J. D. Hamilton, "A new approach to the economic analysis of nonstationary time series and the business cycle," *Econometrica*, vol. 57, no. 2, pp. 357–384, 1989.

[25] M. Karmakar, "Modeling conditional volatility of the Indian stock markets," *Vikalpa*, vol. 30, no. 3, pp. 21–37, 2005.

[26] A. Srivastava, "Volatility of Indian stock market: an empirical evidence," *Asia-Pacific Journal of Management Research and Innovation*, vol. 4, no. 4, pp. 53–61, 2008.

[27] A. L. Eisfeldt, G. Schubert, and M. B. Zhang, "Generative AI and firm values," National Bureau of Economic Research, Cambridge, MA, Working Paper 31222, 2023.

[28] A. L. Eisfeldt and G. Schubert, "AI and finance," National Bureau of Economic Research, Cambridge, MA, Working Paper 33076, 2024.

[29] T. Eloundou, S. Manning, P. Mishkin, and D. Rock, "GPTs are GPTs: labor market impact potential of LLMs," *Science*, vol. 384, no. 6702, pp. 1306–1308, 2024.

[30] E. Brynjolfsson, D. Li, and L. R. Raymond, "Generative AI at work," *The Quarterly Journal of Economics*, vol. 140, no. 2, pp. 889–942, 2025.

[31] S. Noy and W. Zhang, "Experimental evidence on the productivity effects of generative artificial intelligence," *Science*, vol. 381, no. 6654, pp. 187–192, 2023.

[32] S. Peng, E. Kalliamvakou, P. Cihon, and M. Demirer, "The impact of AI on developer productivity: evidence from GitHub

Copilot,” arXiv:2302.06590, 2023.
[33] D. Acemoglu, “The simple macroeconomics of AI,” *Economic Policy*, vol. 40, no. 121, pp. 13–58, 2025.
[34] R. Cerqueti, V. Ficcadenti, and G. Rotundo, “A trillion dollars race: how ChatGPT affects stock prices,” *Future Business Journal*, vol. 11, art. 118, 2025.
[35] Z. Qian, J. Peng, and J. Li, “The impact of generative AI announcements on suppliers: evidence from the stock market,” *Production and Operations Management*, 2025, doi:10.1177/10591478251398333.
[36] “ChatGPT and stock returns,” *Journal of Economics and Finance*, 2026, doi:10.1007/s12197-026-09767-9.
[37] Z. Jiang, D. Xu, and J. Liang, “A deep reinforcement learning framework for the financial portfolio management problem,” arXiv:1706.10059, 2017.
[38] R. J. Shiller, *Narrative Economics: How Stories Go Viral and Drive Major Economic Events*. Princeton, NJ: Princeton University Press, 2019.
[39] NASSCOM, *Technology Sector in India: Strategic Review 2026*. New Delhi, India: National Association of Software and Service Companies, 2026. [Online]. Available: https://nasscom.in/knowledge-center/publications/technology-sector-india-strategic-review-2026
[40] ICICI Direct Research, “Why are IT stocks falling? AI disruption concerns and the IT sector outlook 2026,” Mumbai, India, Equity Research Note, Feb. 2026. [Online]. Available: https://www.icicidirect.com/research/equity/finace/why-are-it-stocks-falling-ai-disruption-concerns-it-sector-outlook-2026
[41] PL Capital Research, “NIFTY IT falls 25% amid AI fears: is a demand cycle repositioning underway?” Mumbai, India, Sector Note, Apr. 2026. [Online]. Available: https://www.plindia.com/news/nifty-it-falls-25-amid-ai-fears-is-a-demand-cycle-repositioning-underway/
[42] Sharpely Research, “Indian IT stocks 2026: AI disruption, price crash, and the midcap opportunity,” Market Analysis, Jul. 2026. [Online]. Available: https://sharpely.in/blogs/indian-stocks-2026-ai-disruption-price-crash-midcaps-hidden-opportunity/
[43] Multibagg Market Pulse, “Nifty IT slides 43% from peak as AI fear hits 2026,” Market Report, Jul. 2026. [Online]. Available: https://www.multibagg.ai/market-pulse/articles/nifty-it-slump-ai-disruption-cmr0snfsz6wbso20jaz1uyej4
[44] Bajaj Broking Research, “Why is Nifty IT falling? AI fears and global tech selloff weigh,” Market Commentary, Jun. 2026. [Online]. Available: https://www.bajajbroking.in/share-market-news/Why-Is-Nifty-IT-Falling-AI-Fears-and-Global-Tech-Selloff-Weigh