Published by
NIGERIAN SOCIETY OF PHYSICAL SCIENCES
Available online @ https://asr.nsps.org.ng/index.php/asr

African Scientific Reports 5 (2026) 618

African Scientific Reports

Stochastic optimization models for energy-resilient supply chains: mitigating cost shocks in Nigerian MSMEs using scientific machine learning

Samuel O. Essang [a,*], Runyi Francis [b]

[a] Mathematics Programme, Department of Physical Sciences, Landmark University, Omu-Aran, Kwara State, Nigeria
[b] Department of Statistics, Federal Polytechnic Ugep, Cross River State, Nigeria

Abstract

Nigerian micro, small and medium-sized enterprises (MSMEs) absorb energy cost shocks with almost no financial buffer. This paper develops a stochastic optimization framework that links published national energy prices to enterprise investment decisions. Using twenty-nine months of official pump price data covering January 2024 to May 2026, we compare four candidate price processes by maximum likelihood estimation. A Merton jump-diffusion process is selected by the Akaike, corrected Akaike and Bayesian information criteria, and the jump term is supported by a boundary-aware parametric bootstrap, whereas mean reversion is not detected, so that for this sample price shocks appear largely permanent rather than temporary. The diffusive drift is statistically indistinguishable from zero, and the jump component accounts for essentially the whole of the estimated log-price drift of 0.373 per year. We embed the fitted process in a two-stage stochastic program with a mean and conditional value-at-risk objective and an energy-balance merit-order recourse, and we add two scientific machine learning components, namely a grey-box drift model and a deep recourse operator. Optimal solar and storage investment reduces the expected annual energy cost of a representative enterprise by 42.8 per cent and the conditional value-at-risk of its operating cost by 86.1 per cent. The shadow price of investment capital reaches about 2.67 naira per naira invested, which shows that, within the model, finance rather than technology or tariff policy is the binding constraint. The learned recourse operator reproduces expected cost to within about one per cent and conditional value-at-risk to within about two and a half per cent, while running about one hundred times faster.

DOI: 10.46481/asr.2026.5.3.618

Keywords: Stochastic optimization, Energy resilience, Jump-diffusion, Scientific machine learning

Article History:
Received: 02 August 2026
Received in revised form: 04 September 2026
Accepted for publication: 11 September 2026
Available online: 22 September 2026

1. Introduction

Micro, small and medium enterprises (MSMEs) form the operating core of the Nigerian economy. The most recent national survey conducted jointly by the Small and Medium Enterprises Development Agency of Nigeria and the National Bureau of Statistics counted 39,654,385 such enterprises [1]. They account for 96.9 per cent of all businesses, 46.31 per cent of gross domestic product

*Corresponding Author Tel. No.: +234-806-623-6521.
Email address: essang.samuel@lmu.edu.ng (Samuel O. Essang)

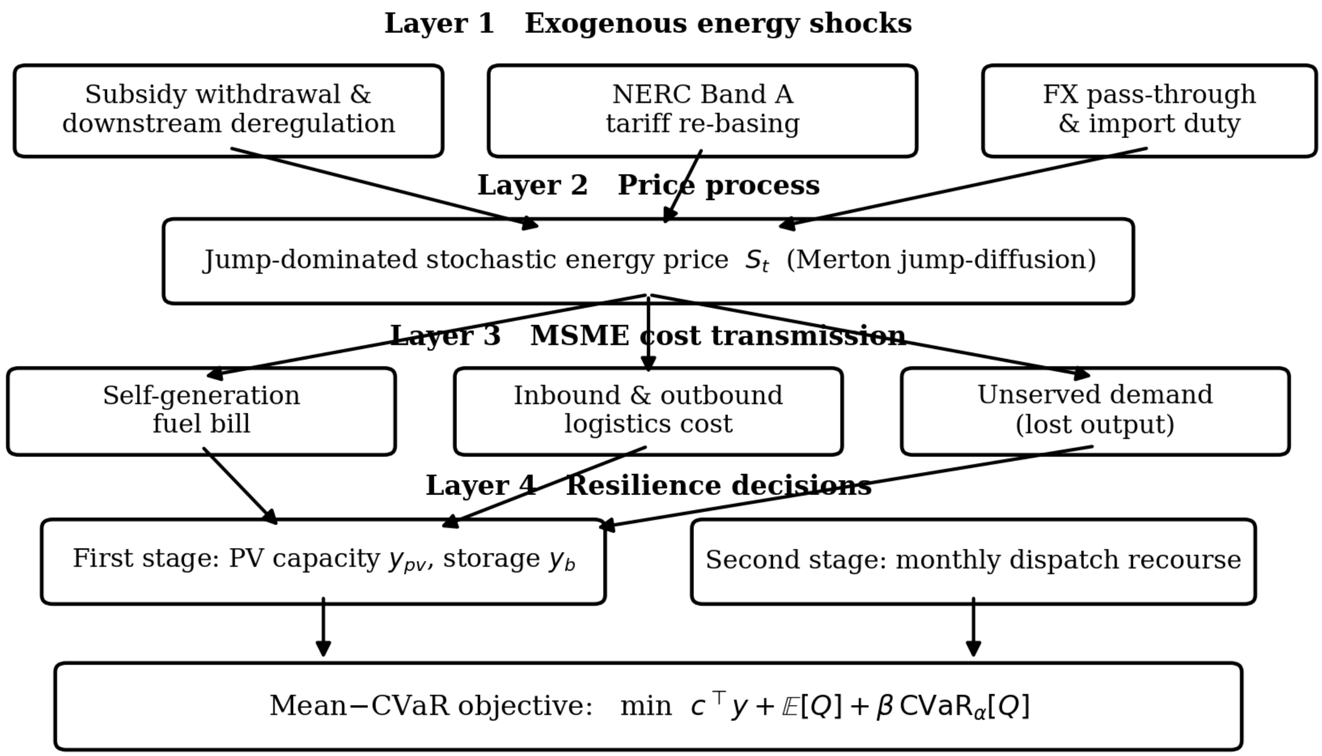

Figure 1: Conceptual framework linking exogenous energy policy shocks to enterprise resilience decisions. Layer 1 lists the shock sources; Layer 2 represents them as a jump-dominated stochastic price; Layer 3 shows the three channels of cost transmission to the enterprise; Layer 4 sets out the two-stage decision structure and the mean-CVaR objective.

and 87.9 per cent of employment, yet contribute only 6.21 per cent of the export basket [1]. This combination matters: a cost shock reaching Nigerian MSMEs reaches almost the entire domestic productive base at once, and the firms absorbing the shock earn little foreign exchange with which to hedge it.

Energy is the channel through which such shocks arrive most forcefully. Nigerian firms do not consume energy at the price posted by the public grid, because for most of them the public grid is not the marginal supplier. The African Development Bank reports that 70.7 per cent of firms in Nigeria own or share a generator, against a global average of about 30.5 per cent, and that electricity outages cost Nigerian firms about 3 per cent of annual sales [2, 3]. A firm in this position pays a blended energy price: part of its demand is met from the grid at the regulated tariff, and the remainder from a petrol or diesel generator at a price set by the downstream fuel market. The second component dominates both the level and the variability of the bill.

Two policy events reshaped that blended price. In May 2023 the Federal Government withdrew the petrol subsidy and deregulated the downstream petroleum sector. In April 2024 the Nigerian Electricity Regulatory Commission re-based the Band A end-user tariff from about 66-68 naira per kilowatt-hour to 225 naira per kilowatt-hour, a rise of roughly 240 per cent for customers guaranteed at least twenty hours of daily supply [4]. Neither change was a gradual adjustment; both were discrete steps, and both were followed by further discrete steps as refinery entry, foreign-exchange movements and import-duty decisions worked through the market. Figure 1 summarises how these exogenous shocks are transmitted to the enterprise and turned into an investment decision.

This is the empirical feature that motivates the present paper. Standard energy-planning models represent price uncertainty as a diffusion, in which the price wanders continuously and, in mean-reverting formulations, is pulled back toward a long-run level. Under that view a cost shock is a temporary inconvenience and the rational response is to wait. If instead the price process is dominated by discrete jumps with no reversion, a shock is largely permanent, waiting destroys value, and the rational response is to restructure the firm's energy supply. Which description is the better approximation is an empirical question that has not, to our knowledge, been examined for the post-subsidy Nigerian market using formal model selection. Examining it is the first contribution of this paper.

The second contribution is decision-theoretic. Once the price process is identified, the enterprise faces a problem with a natural two-stage structure. Before the year begins it chooses durable capacity: how many kilowatts-peak of solar photovoltaic capacity to install and how many kilowatt-hours of battery storage to pair with it. These are first-stage decisions, taken under uncertainty and not reversible month by month. Once each month arrives and the fuel price and grid availability are observed, the firm chooses how to meet its load from the available sources. These are second-stage recourse decisions. We formulate this as a two-stage stochastic program and, because a small firm cares about survival and not only about averages, we optimise a mean-conditional-value-at-risk objective rather than expected cost alone.

The third contribution concerns method. Scientific machine learning (SciML) integrates mechanistic structure with data-driven components, and has advanced rapidly in building energy systems [5] and industrial process modelling [6]. Its promise for energy economics is clear, but so is its risk: the Nigerian pump-price record contains twenty-nine monthly observations, far too few to train an unconstrained network. We therefore apply SciML in two distinct roles and report both outcomes honestly. In the first role, learning the price drift, we test a grey-box universal differential equation and a black-box network against the fitted parametric process on held-out data. In the second role, approximating the recourse function of the optimization problem, we train a deep operator surrogate. The two roles produce opposite verdicts, and we regard the negative result as being as informative as the positive one.

The remainder of the paper is organised as follows. Section 2 reviews the relevant literature. Section 3 describes the data, develops the price model and the stochastic program, and presents the scientific machine learning layer. Section 4 reports results. Section 5 discusses them, states limitations, and draws policy implications. Section 6 concludes.

1.1. Novelty of the study

The novel elements of this study can be stated as six specific contributions.

(1) *A formal comparison of candidate price processes for the Nigerian post-subsidy pump price.* We fit four nested and non-nested specifications by maximum likelihood to the same data and rank them by the Akaike information criterion, its small-sample correction, and the Bayesian information criterion, supported by a likelihood-ratio test whose reference distribution is obtained by parametric bootstrap. Within this comparison, adding jumps to a geometric Brownian motion improves the fit decisively, whereas adding mean reversion does not improve it, so that the data are best described by a jump-dominated specification with no detectable mean reversion.

(2) *Attribution of the estimated drift to discrete events.* In the selected model the diffusive drift is statistically indistinguishable from zero, while the jump component accounts for essentially all of the estimated log-price drift of 0.373 per year. This decomposition converts a familiar macroeconomic narrative into an estimated quantity.

(3) *A two-stage mean-CVaR resilience model driven by a jump process with an explicit energy-balance recourse.* The second stage is not a reduced-form cost curve. It is an explicit merit-order dispatch in which solar energy serves daytime load directly but serves night-time load only through the battery, subject to depth of discharge, round-trip efficiency and a monthly cycle limit.

(4) *Identification of capital, within the model, as the binding constraint.* By sweeping the investment budget we recover the shadow price of investment capital, which reaches about 2.67 naira of annual objective improvement per naira invested at low budgets and falls to zero beyond about seven million naira. This locates, within the model, where a relaxation of the budget would be most valuable.

(5) *A separation of the roles in which scientific machine learning helps and harms.* For price drift, with twenty-eight transitions, the black-box network attains the best in-sample fit and the worst out-of-sample error, and the parametric process generalises best. For the recourse operator, where training data can be generated without limit, the neural surrogate is highly accurate and about two orders of magnitude faster. We report both.

(6) *Storage identified as a risk instrument rather than a least-cost energy asset.* Under a risk-neutral objective the optimal battery is essentially zero. Storage enters the solution only as the risk-aversion weight rises, which has a direct implication for how storage subsidies should be justified.

2. Literature review

2.1. Stochastic modelling of energy prices

Energy prices are widely recognised as poorly described by Gaussian models. Ganczarek-Gamrot and co-authors examine seven European bidding zones over 2023-2024 and document strongly skewed, heavy-tailed price distributions with pronounced outliers, concluding that non-parametric methods are preferable and that forecast quality deteriorates as volatility rises [7]. Ignatieva and Wong compare affine and non-affine jump-diffusion models for crude oil and find that jump components are necessary to reproduce observed dynamics [8]. Ben Romdhane and Boubaker combine heterogeneous autoregressive structure with a recurrent network and a generalised autoregressive conditional heteroskedasticity component for crude-oil futures volatility [9]. These studies establish that jumps matter in mature, liquid markets. Our setting differs in kind: the Nigerian pump price is driven not primarily by traded market microstructure but by administrative and structural events such as subsidy withdrawal, refinery entry and duty changes, and whether the resulting series is best described by a jump process, and whether it reverts, has not been established by formal model comparison.

2.2. Stochastic optimization for resilient supply chains and energy systems

The two-stage stochastic program is the standard formalism for investment under uncertainty with operational recourse. Moghadaspoor and co-authors design an energy-resilient closed-loop supply chain managed by a third-party logistics provider, combining multi-objective optimization with data envelopment analysis [10]; Mozaffari and colleagues develop a two-stage framework for distribution-network resilience under extreme events with electric-vehicle integration [11]; Oszczypala and co-authors combine continuous-time Markov chains with Monte Carlo simulation for redundancy optimization [12]; and Jadhav and Kumar couple a continuous-time Markov chain with particle-swarm optimization to raise the availability of an agricultural wireless sensor network [13]; and, on the design side, Şenol and Murat develop a sequential solution heuristic for continuous facility-layout optimization [14]. Within the Nigerian power-system context specifically, Adepoju and colleagues apply a hybrid particle-swarm and bat-algorithm metaheuristic to the economic dispatch of the national thermal generation fleet, minimising generation cost subject to the system's operating constraints [15]. Risk-averse formulations have become standard where tail outcomes matter: Wang, Liu and Bai design resilient supply-chain networks under partial distribution information [16]; Wang and Zhong develop a distributionally robust approach to humanitarian network design under correlated disruptions [17]; Davoodi, Gholamian and Hanne incorporate conditional risk in a multi-level, multi-product supply chain [18]; and Chakrabarty and colleagues address time-based redeployment for reliable coverage [19]. The gap we address is that these models are almost always calibrated to firms with balance sheets large enough to absorb a bad year, whereas the Nigerian micro-enterprise has no such buffer and an unusually narrow recourse set: it cannot renegotiate supply contracts, hedge on a futures market or relocate production.

2.3. *Energy cost shocks and MSME resilience*

A growing empirical literature documents the consequences of Nigeria's energy reforms for small firms. Gumel examines the joint impact of fuel-subsidy removal and foreign-exchange harmonisation on Nigerian small and medium enterprises and reports systemic cost escalation [20]. Priya and Mathew review business-model innovation as a driver of MSME sustainability and resilience, identifying resource constraints and market uncertainty as recurring barriers [21]. Achmad and Subriadi study blockchain adoption among MSMEs and find that inadequate infrastructure, skills and capital obstruct technology adoption [22], a finding that anticipates our own conclusion about capital rationing. Edoja and colleagues analyse macroeconomic drivers of price volatility in a Nigerian agricultural value chain [23], and Edobor and Sambo-Magaji consider small enterprises within sustainable development [24]. The renewable-resource base that on-site investment would draw on is itself increasingly well documented for Nigeria; Ahile and colleagues, for instance, assess the geothermal potential of the Benue Trough as an alternative energy source using high-resolution aeromagnetic data [25]. This literature is predominantly descriptive or survey-based: it establishes that the shock is severe but does not provide an optimization model that tells a specific firm what to do about it. Abba, Balta-Ozkan and Drew move toward the decision problem by evaluating investor-focused risk in decentralised Nigerian renewable energy [26], and Sim and Griffiths examine the resilience of renewable-energy supply chains [27], but neither addresses enterprise-level capacity choice under a calibrated stochastic price.

2.4. *Scientific machine learning*

Liu and co-authors present scientific machine learning as a paradigm for building technology, integrating physics-based models with data-driven learning to obtain scalable and domain-aware solutions [5], while Mohammadghasemi and colleagues embed neural components inside the physical laws of an industrial primary-separation vessel and then recover an interpretable relation through symbolic regression [6]. Gong and Cheung develop an adaptive physics-informed neural operator for high-dimensional reliability analysis of systems governed by stochastic differential equations [28], and Zhang and colleagues use a physics-informed network as a surrogate inside a many-objective optimization algorithm for coupled multi-energy systems [29]. Two lessons from this literature shape our design. First, embedding known structure reduces the data requirement, which matters acutely when only twenty-nine observations exist. Second, the most reliable industrial payoff from SciML so far has come from surrogate modelling of expensive computations rather than from discovery of unknown dynamics. Our results support both lessons, including where they imply that machine learning should not be used.

2.5. *Digital technology and supply chain resilience*

Wei and co-authors review digital technology and supply-chain resilience and identify predictive analytics and simulation-based stress testing as dominant themes [30], while Li and colleagues combine deep learning, survival analysis and explainable artificial intelligence for end-to-end resilience management [31]. Related reliability-optimization work includes Singh and Singh on dynamic reliability of weighted cold-standby systems [32], Shen, Fouladirad and Grall on solar-farm performance degradation under condition-based maintenance [33], Bhandari, Kumar and Ram on hybrid metaheuristics for redundancy allocation [34], Kumar and colleagues on multistate performance with human error [35], Kumar and Kumar on Markov decision processes for system reliability [36], Patel and co-authors on integrated reliability assessment of industrial tooling [37], and Kumar and colleagues on availability evaluation of solar photovoltaic systems using Markov modelling with a metaheuristic search [38]. These contributions assume an enterprise with digital infrastructure and data-collection capacity, which the Nigerian micro-enterprise typically lacks. Our framework is deliberately built on publicly available national statistics so that it can be applied without any firm-level data system.

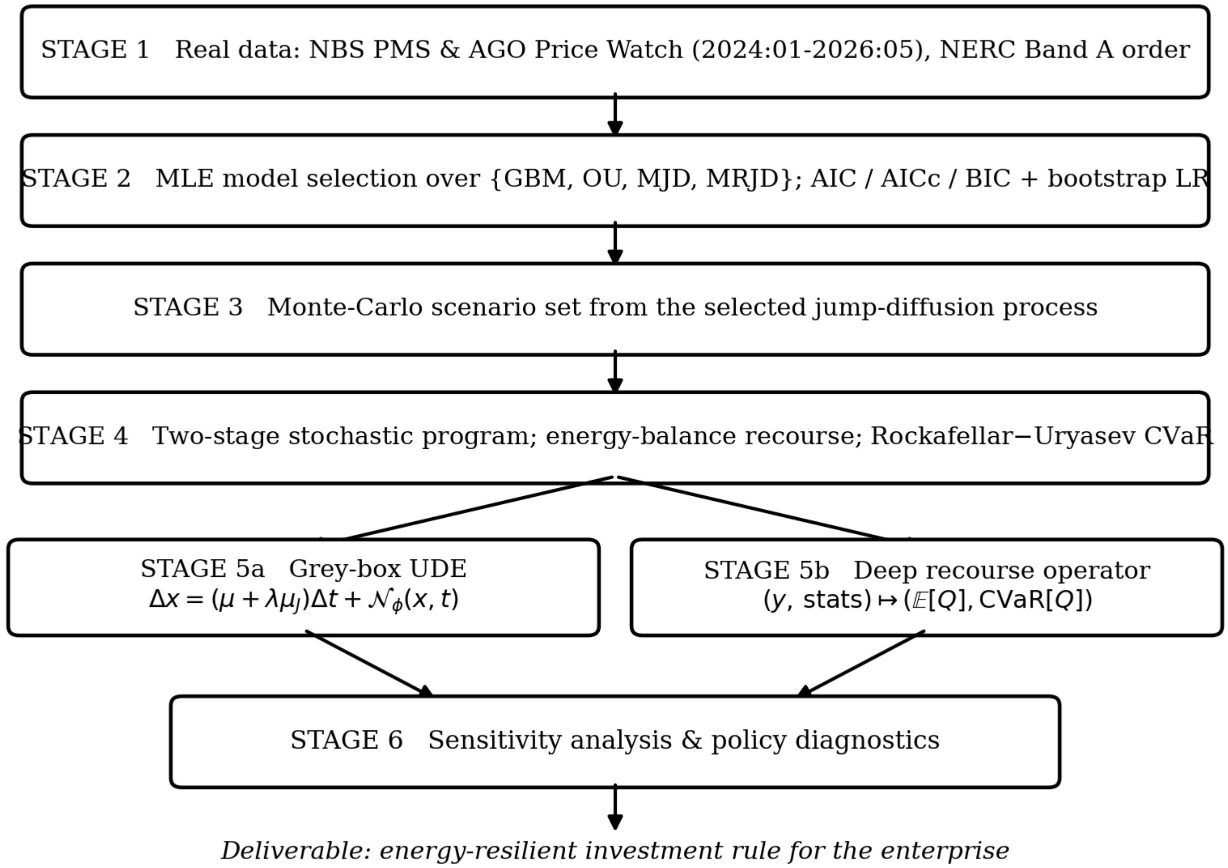

Figure 2: Methodological architecture of the study. Stages 1 to 4 form the mechanistic pipeline from published data to the solved stochastic program, with the parametric jump test assessed by bootstrap and the recourse solved as an energy balance. Stages 5a and 5b are the two scientific machine learning components, applied respectively to drift learning and to recourse-function approximation.

3. Materials and methods

3.1. Data sources and description

3.1.1. Sources

All price data in this study are published official statistics; we use no simulated, imputed or proprietary price observations. The primary series is the national average retail price of Premium Motor Spirit (PMS, petrol) published monthly by the National Bureau of Statistics in the *Premium Motor Spirit (Petrol) Price Watch* [39]. The Bureau collects these prices from sample retail outlets across all 774 local government areas in the thirty-six states and the Federal Capital Territory, from more than 10,000 respondents, using over 700 field officers under quality control. We use the complete monthly series from January 2024 to May 2026, giving twenty-nine observations; the full series is reproduced in Appendix A.

The secondary series is the national average retail price of Automotive Gas Oil (AGO, diesel) from the companion *Automotive Gas Oil (Diesel) Price Watch* [40]. Coverage of published monthly figures is incomplete for our window, so we use the seventeen months we were able to verify, and only for descriptive comparison, never for estimation. Electricity tariffs are taken from the Nigerian Electricity Regulatory Commission Band A orders under the Service-Based Tariff regime [4]; the Band A rate moved from approximately 66-68 naira per kilowatt-hour to 225 naira per kilowatt-hour with effect from 3 April 2024, was subsequently revised to 206.80 and then 209.50 naira per kilowatt-hour, and we adopt 209.50 naira per kilowatt-hour as the base case. Structural parameters of the enterprise sector are taken from the SMEDAN/NBS National MSME Survey [1] and the African Development Bank's *African Economic Outlook 2026* [2], supplemented by the World Bank Enterprise Surveys for the international comparison of generator ownership [3]. Figure 2 summarises the full methodological pipeline from these sources to the resilience decision.

Table 1: Descriptive statistics of the NBS energy price series, January 2024 to May 2026.

Statistic	PMS (petrol)	AGO (diesel)	PMS log-returns
Observations	29	17	28
Mean	1041.57	1457.41	0.0311
Standard deviation	242.73	159.18	0.0811
Minimum	668.30	1257.06	-0.1872
Maximum	1596.25	1813.81	0.2158
Coefficient of variation	0.233	0.109	-
Skewness	-	-	0.312
Kurtosis	-	-	4.397
Jarque-Bera statistic	-	-	2.729
Jarque-Bera p-value	-	-	0.256
Lag-1 autocorrelation	-	-	0.384

Prices in naira per litre. Kurtosis is on the raw scale, for which the Gaussian value is 3.0.

3.1.2. Descriptive statistics

Table 1 reports summary statistics. Over the twenty-nine months the PMS price averaged 1041.57 naira per litre with a standard deviation of 242.73 naira, a coefficient of variation of 0.233, and a range from 668.30 naira in January 2024 to 1596.25 naira in May 2026; the maximum exceeds the minimum by 138.9 per cent. For a firm whose marginal energy source is a petrol generator, this is the range over which its unit energy cost moved within a little over two years.

The twenty-eight monthly log-returns have a mean of 0.0311 and a standard deviation of 0.0811. The distribution is positively skewed at 0.312 and, more importantly, has a kurtosis of 4.397 against the Gaussian value of 3.0. The monthly log-returns range from -0.187 to +0.216, corresponding to single-month price changes of -17.1 and +24.1 per cent. The first-order autocorrelation of the returns is 0.384, decaying to 0.040 at lag two, indicating short-lived momentum rather than the negative autocorrelation that mean reversion in levels would generate.

We note explicitly that the Jarque-Bera statistic is 2.729 with a p-value of 0.256, so normality is not rejected at conventional levels. With twenty-eight observations that test has very low power against the alternatives of interest, and we do not treat it as evidence for a Gaussian model; the formal comparison in Section 4.2 uses likelihood-based criteria instead, which are better suited to this sample size. Figure 3 plots both series with the principal market events annotated. The visual impression is not one of continuous drift but of extended plateaux interrupted by abrupt steps, which is precisely the pattern a jump process is designed to represent.

3.2. Model formulation

3.2.1. Candidate price processes

Let S_t denote the national average PMS price in naira per litre at month t and let $x_t = \log S_t$. We consider four specifications. The first is a geometric Brownian motion,

$$dx_t = \mu\, dt + \sigma\, dW_t, \tag{1}$$

where W_t is a standard Wiener process and σ has units of per square-root year. The second is an Ornstein-Uhlenbeck process in the log price,

$$dx_t = \kappa\,(\theta - x_t)\, dt + \sigma\, dW_t, \tag{2}$$

in which $\kappa > 0$ is the speed of mean reversion and θ the long-run log price. The third is the Merton jump-diffusion,

$$dx_t = \mu\, dt + \sigma\, dW_t + J_t\, dN_t, \tag{3}$$

where N_t is a Poisson process with intensity λ and the jump sizes satisfy $J_t \sim \mathcal{N}\big(\mu_J, \sigma_J^2\big)$, independently of W_t. The fourth adds the same jump term to the mean-reverting drift,

$$dx_t = \kappa\,(\theta - x_t)\, dt + \sigma\, dW_t + J_t\, dN_t. \tag{4}$$

The one-month transition density is obtained by conditioning on the number of jumps n in the interval $\Delta = 1/12$, which is Poisson with mean $\lambda\Delta$, and summing. Writing $r_t = x_{t+1} - x_t$ and m_t for the conditional diffusive mean, the density is the Poisson mixture

$$f(r_t) = \sum_{n=0}^{K} \frac{e^{-\lambda\Delta}(\lambda\Delta)^n}{n!}\, \phi\Big(r_t;\; m_t + n\mu_J,\; \sigma^2\Delta + n\sigma_J^2\Big), \tag{5}$$

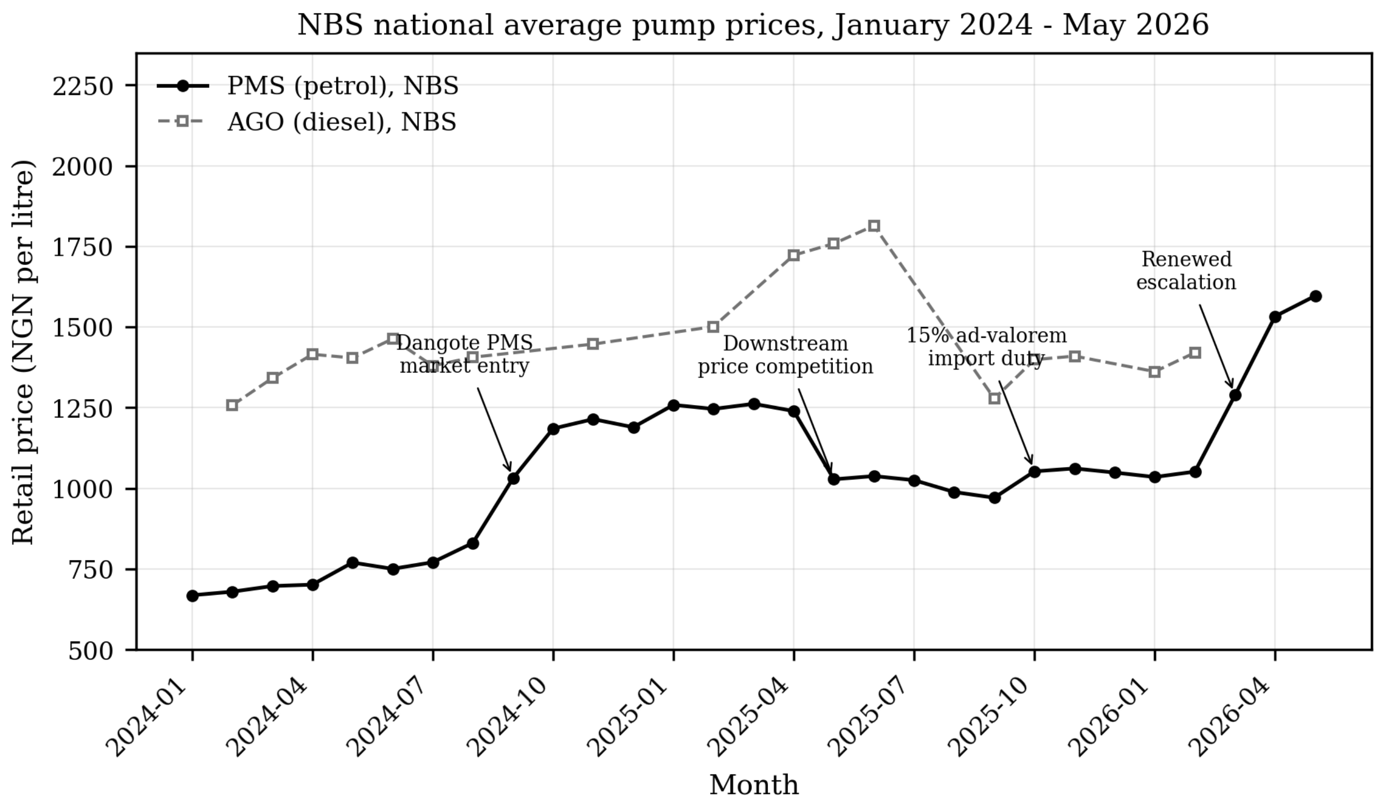

Figure 3: National average retail prices of Premium Motor Spirit and Automotive Gas Oil, January 2024 to May 2026, from the National Bureau of Statistics monthly Price Watch reports. Annotations mark the principal recorded market events. The pattern of extended plateaux interrupted by abrupt steps motivates a jump-process specification.

where $\phi(\cdot\,;a,b)$ is the Gaussian density with mean a and variance b, $m_t = \mu\Delta$ for the non-reverting specifications and $m_t = \kappa(\theta - x_t)\Delta$ for the reverting ones, and the truncation $K = 7$ is large enough that the neglected Poisson mass is below 1.4×10^{-7} at the fitted intensity. Using the full Poisson mixture, rather than a one-jump approximation, ensures that the estimation model coincides with the compound-Poisson process used to generate scenarios in Section 3.2.4. For the non-reverting specifications (M1 and M3) the diffusion transition is exact. For the reverting specifications (M2 and M4) we use the exact Ornstein-Uhlenbeck transition, with conditional mean $(\theta - x_t)(1 - e^{-\kappa\Delta})$ and variance $\sigma^2(1 - e^{-2\kappa\Delta})/(2\kappa)$ in place of the Euler forms; the only residual approximation is that, under mean reversion, the compound-Poisson jump term is convolved with this Gaussian as though the jumps arrive at the interval boundary rather than being propagated by the reverting drift within the month. Since the selected model is non-reverting, this approximation does not affect the reported estimates for M3. The log-likelihood is

$$\ell(\Theta) = \sum_{t=1}^{n} \log f(r_t), \tag{6}$$

maximised by multi-start Nelder-Mead followed by a Powell polish, with sixty random restarts per specification. Variance parameters are estimated in logarithms to enforce positivity, and reported standard errors are obtained from the numerical observed information matrix with the delta method applied to the transformed parameters.

3.2.2. Model selection criteria

Because the specifications differ in dimension we rank them by the Akaike information criterion, its small-sample correction, and the Bayesian information criterion,

$$\mathrm{AIC} = 2k - 2\ell, \quad \mathrm{AIC_c} = \mathrm{AIC} + \frac{2k(k+1)}{n-k-1}, \quad \mathrm{BIC} = k\log n - 2\ell, \tag{7}$$

where k is the number of free parameters and n the number of returns. The corrected criterion matters here: with $n = 28$ the ratio n/k is small for the richer models, and an uncorrected criterion would favour over-parameterisation. Two of the six model pairs are regular nested comparisons in which jumps are added to a diffusion, namely M1 within M3 and M2 within M4, obtained by setting $\lambda = 0$. For these we report the likelihood-ratio statistic

$$\mathrm{LR} = 2\left(\ell_{\text{full}} - \ell_{\text{restricted}}\right). \tag{8}$$

Because $\lambda = 0$ lies on the boundary of the parameter space and the jump-size parameters (μ_J, σ_J) are unidentified under the null, the usual χ^2 reference distribution for (8) is not valid; we therefore obtain the p-value of the jump test by parametric bootstrap, simulating under the fitted no-jump model and refitting both specifications on each replicate. The mean-reversion comparisons (M1 against M2, and M3 against M4) are not regular nested pairs under the drift parameterisation of equations (1)-(4), since setting $\kappa = 0$ removes the drift rather than recovering an arbitrary constant μ; these are therefore assessed by the information criteria only, not by a likelihood-ratio test.

3.2.3. Jump identification

Having fitted the selected model we recover, for each observed month, the posterior probability that the return contained at least one jump. By Bayes' rule, with $p_0 = e^{-\lambda\Delta}\,\phi(r_t; m_t, \sigma^2\Delta)$ the no-jump component and $f(r_t)$ the full mixture density of equation (5),

$$P(\text{at least one jump} \mid r_t) = 1 - \frac{p_0}{f(r_t)}. \tag{9}$$

This quantity is a diagnostic rather than an estimate: it allows the statistically identified jump months to be compared against the documented policy and market record, providing an external validity check on the fitted process.

3.2.4. Scenario generation

Scenarios are generated by Euler discretisation of the selected process with an exact compound-Poisson jump overlay, matching the estimation density of equation (5). For path ω and month t,

$$x_t^{(\omega)} = x_{t-1}^{(\omega)} + \mu\Delta + \sigma\sqrt{\Delta}\,z_t^{(\omega)} + \sum_{i=1}^{N_t^{(\omega)}} J_{t,i}^{(\omega)}, \quad N_t^{(\omega)} \sim \mathrm{Poisson}(\lambda\Delta), \tag{10}$$

with $z_t^{(\omega)} \sim \mathcal{N}(0,1)$ and $J_{t,i}^{(\omega)} \sim \mathcal{N}(\mu_J, \sigma_J^2)$, all mutually independent, and $S_t^{(\omega)} = \exp(x_t^{(\omega)})$. We generate $\Omega = 4000$ paths over a twelve-month horizon from the last observed price, with a fixed random seed so that every reported figure is reproducible.

Grid availability is modelled as a separate stochastic process. Let $a_t^{(\omega)} \in [0,1]$ be the share of monthly demand the public grid can serve. In a normal month $a_t^{(\omega)} \sim \mathrm{Beta}(mc, (1-m)c)$ with mean $m = 0.55$ and concentration $c = 6$, and with probability 0.10 the month is a severe-failure month in which the mean falls to 0.12. The realised sample mean grid share is 0.51, consistent with the evidence that Nigerian firms self-generate roughly half of their consumption [2, 3]. Fuel price and grid availability are modelled as independent; this assumption, and its likely direction of bias, is discussed in Section 5.3.

3.2.5. The two-stage stochastic program

The enterprise chooses first-stage capacities $y = (y_{pv}, y_b)$, where y_{pv} is installed photovoltaic capacity in kilowatts-peak and y_b is battery capacity in kilowatt-hours. Writing $c = (c_{pv}, c_b)$ for unit installed costs, the annualised capital charge is

$$C(y) = c_{pv}y_{pv}\,\mathrm{CRF}(\rho, L_{pv}) + c_b y_b\,\mathrm{CRF}(\rho, L_b) + \gamma\left(c_{pv}y_{pv} + c_b y_b\right), \tag{11}$$

where γ is the annual operation-and-maintenance rate, L denotes asset life, and the capital recovery factor is

$$\mathrm{CRF}(\rho, L) = \frac{\rho(1+\rho)^L}{(1+\rho)^L - 1}. \tag{12}$$

The full problem is

$$\min_{y\in\mathcal{Y}} \; C(y) + \mathbb{E}\left[Q(y,\omega)\right] + \beta\,\mathrm{CVaR}_\alpha\left[Q(y,\omega)\right], \tag{13}$$

subject to the investment budget constraint

$$\mathcal{Y} = \left\{(y_{pv}, y_b) \in \mathbb{R}_+^2 \;:\; c_{pv}y_{pv} + c_b y_b \le B\right\}. \tag{14}$$

Here $Q(y,\omega)$ is the second-stage annual operating cost under scenario ω, α is the confidence level and $\beta \ge 0$ the risk-aversion weight; setting $\beta = 0$ recovers the risk-neutral problem, and increasing β traces the cost-risk efficient frontier. Following Rockafellar and Uryasev, conditional value-at-risk admits the empirical representation

$$\mathrm{CVaR}_\alpha\,[Q] = \min_{\eta\in\mathbb{R}} \left\{ \eta + \frac{1}{(1-\alpha)\Omega} \sum_{\omega=1}^{\Omega} [Q(y,\omega) - \eta]^+ \right\}, \tag{15}$$

where $[\cdot]^+ = \max(\cdot, 0)$. We evaluate it directly from the sorted scenario costs, taking the value-at-risk as the $\lceil \alpha\Omega \rceil$-th order statistic and the conditional value-at-risk as the mean of the tail at and beyond it.

3.2.6. Energy-balance recourse

The second stage is an explicit dispatch problem rather than a fitted cost curve, but it is a monthly energy-balance model and not an hour-by-hour simulation; the consequences of this aggregation are stated in Section 5.3. Let E be monthly electricity demand in kilowatt-hours, τ the Band A tariff, ϱ the specific fuel consumption of the generator in litres per kilowatt-hour, and v the value of lost load. Monthly photovoltaic generation per unit capacity is

$$g_{pv} = H\,\mathrm{PR}\,d, \tag{16}$$

where H is peak sun hours per day, PR the performance ratio and d the mean days per month, so total generation is $G = g_{pv}\,y_{pv}$. Solar output is available only in daylight, so with φ the share of monthly demand coincident with photovoltaic output the direct solar service is

$$s_{\mathrm{dir}} = \min(G,\ \varphi E)\,. \tag{17}$$

Any surplus can serve night-time load only through the battery, whose monthly deliverable energy is limited by usable depth of discharge δ, round-trip efficiency η and the number of cycles ν,

$$s_{\mathrm{sto}} = \min(G - s_{\mathrm{dir}},\quad y_b\,\delta\,\eta\,\nu,\quad (1-\varphi)E)\,, \tag{18}$$

with total solar service $s = s_{\mathrm{dir}} + s_{\mathrm{sto}}$. Equation (18) is an energy-balance limit: it caps the energy the battery can deliver over the month but does not track a state of charge, charge and discharge power limits, or the chronological matching of surplus generation to night-time demand within the month. It nonetheless prevents the physically impossible outcome that a large array with no storage supplies the firm at night. Dispatch of the residual demand follows merit order, since marginal costs satisfy $0 < \tau < \varrho S_t < v$ over the relevant price range. Residual demand after solar is $R_t^{(\omega)} = E - \min(s, E)$; grid service is

$$u_t^{(\omega)} = \min\left(a_t^{(\omega)} E,\ R_t^{(\omega)}\right), \tag{19}$$

which likewise treats the monthly grid-available energy as fully allocatable to residual demand; generator service is capped by installed genset capability χE,

$$h_t^{(\omega)} = \min\left(R_t^{(\omega)} - u_t^{(\omega)},\ \chi E\right), \tag{20}$$

and any remaining demand is unserved,

$$w_t^{(\omega)} = R_t^{(\omega)} - u_t^{(\omega)} - h_t^{(\omega)}. \tag{21}$$

The annual operating cost under scenario ω is therefore

$$Q(y,\omega) = \sum_{t=1}^{T} \left[\tau\, u_t^{(\omega)} + \varrho\, S_t^{(\omega)}\, h_t^{(\omega)} + v\, w_t^{(\omega)} \right]. \tag{22}$$

Because $Q(y, \omega)$ is evaluated in closed form for every scenario, the first-stage problem is solved by exhaustive search over a lattice in (y_{pv}, y_b) with a resolution of 0.25 kilowatts-peak and 0.5 kilowatt-hours, discarding infeasible points; this returns a global optimum on the lattice without any convexity assumption.

3.2.7. Parameter values

Table 2 lists the parameters of the representative enterprise, a small agro-processing firm consuming 600 kilowatt-hours per month, a scale consistent with the micro and small categories that dominate the Nigerian distribution [1]. Peak sun hours of 4.9 per day reflect the Nigerian annual mean; the performance ratio of 0.78 is standard for fixed-tilt installations in hot climates; capital costs of 900,000 naira per kilowatt-peak and 450,000 naira per kilowatt-hour reflect installed prices including balance of system and labour; and the 25 per cent cost of capital reflects Nigerian commercial lending conditions. Every parameter is subjected to sensitivity analysis in Section 4.7.

Table 2: Parameters of the representative enterprise.

Symbol	Parameter	Value	Unit
E	Monthly electricity demand	600	kWh/month
T	Planning horizon	12	months
τ	NERC Band A tariff	209.50	NGN/kWh
ϱ	Generator specific fuel consumption	0.42	litres/kWh
χ	Generator capability cap	0.80	share of demand
v	Value of lost load	1500	NGN/kWh
H	Peak sun hours	4.9	hours/day
PR	Photovoltaic performance ratio	0.78	-
d	Mean days per month	30.44	days
φ	Daytime demand coincidence share	0.55	-
η	Battery round-trip efficiency	0.90	-
δ	Usable depth of discharge	0.85	-
ν	Battery cycles per month	30	cycles
c_{pv}	Installed photovoltaic cost	900,000	NGN/kWp
c_b	Installed battery cost	450,000	NGN/kWh
L_{pv}	Photovoltaic asset life	10	years
L_b	Battery asset life	8	years
ρ	Cost of capital	0.25	per year
γ	Annual operation and maintenance rate	0.02	of capital cost
B	Investment budget (base case)	6,000,000	NGN
α	CVaR confidence level	0.95	-
Ω	Monte-Carlo scenarios	4000	paths

3.3. *Scientific machine learning layer*

3.3.1. *Grey-box drift learning*

The first SciML component tests whether a learned correction improves the drift of the selected model. We retain the deterministic drift implied by the fitted jump-diffusion and add a neural residual, giving a universal differential equation

$$\Delta x_t = \underbrace{(\mu + \lambda\mu_J)\,\Delta}_{\text{selected-model drift}} + \underbrace{\mathcal{N}_\phi(x_t, \varsigma_t)}_{\text{learned residual}} + \varepsilon_t, \tag{23}$$

where the mechanistic term is the compensated drift of the selected Merton model, not the rejected mean-reverting term, ς_t is normalised calendar time and $\mathcal{N}_\phi$ is a feed-forward network with two hidden layers of eight hyperbolic-tangent units. The mechanistic parameters are fixed at their calibrated values, so the network is responsible only for what the selected model leaves unexplained. The training objective is the penalised mean squared residual

$$\mathcal{L}(\phi) = \frac{1}{n_{\mathrm{tr}}} \sum_{t=1}^{n_{\mathrm{tr}}} \Big[\Delta x_t - (\mu + \lambda\mu_J)\Delta - \mathcal{N}_\phi(x_t, \varsigma_t)\Big]^2 + \varpi\|\phi\|_2^2, \tag{24}$$

with weight decay $\varpi = 10^{-3}$, minimised by Adam over 4000 iterations. For comparison we train an unconstrained network of identical architecture that predicts Δx_t directly with no mechanistic term. Both are trained on the first twenty-two transitions and evaluated on the final six, which are never seen during training.

3.3.2. *Deep recourse operator*

The second SciML component addresses a computational rather than a statistical problem. Evaluating $\mathbb{E}[Q]$ and $\mathrm{CVaR}_\alpha[Q]$ requires dispatching 4000 scenarios over twelve months for each candidate (y_{pv}, y_b), so re-solving the program as prices update is expensive. We train a surrogate

$$\mathcal{S}_\psi : \Big(y_{pv},\, y_b,\, \bar{S},\, \hat{\sigma}_S,\, \hat{\varsigma}_S,\, S_{90},\, \bar{a},\, \hat{\sigma}_a\Big) \longmapsto \left(\mathbb{E}[Q],\ \mathrm{CVaR}_\alpha[Q]\right), \tag{25}$$

whose inputs summarise the price scenario set by its mean $\bar{S}$, standard deviation $\hat{\sigma}_S$, skewness $\hat{\varsigma}_S$ and ninetieth percentile S_{90}, and the grid availability by its mean $\bar{a}$ and standard deviation $\hat{\sigma}_a$. The higher moments and the tail quantile are included precisely because the mean and standard deviation alone do not determine the tail of a nonlinear dispatch cost, so a conditional value-at-risk map requires information about distributional shape. These summaries do not uniquely determine the full scenario distribution or its

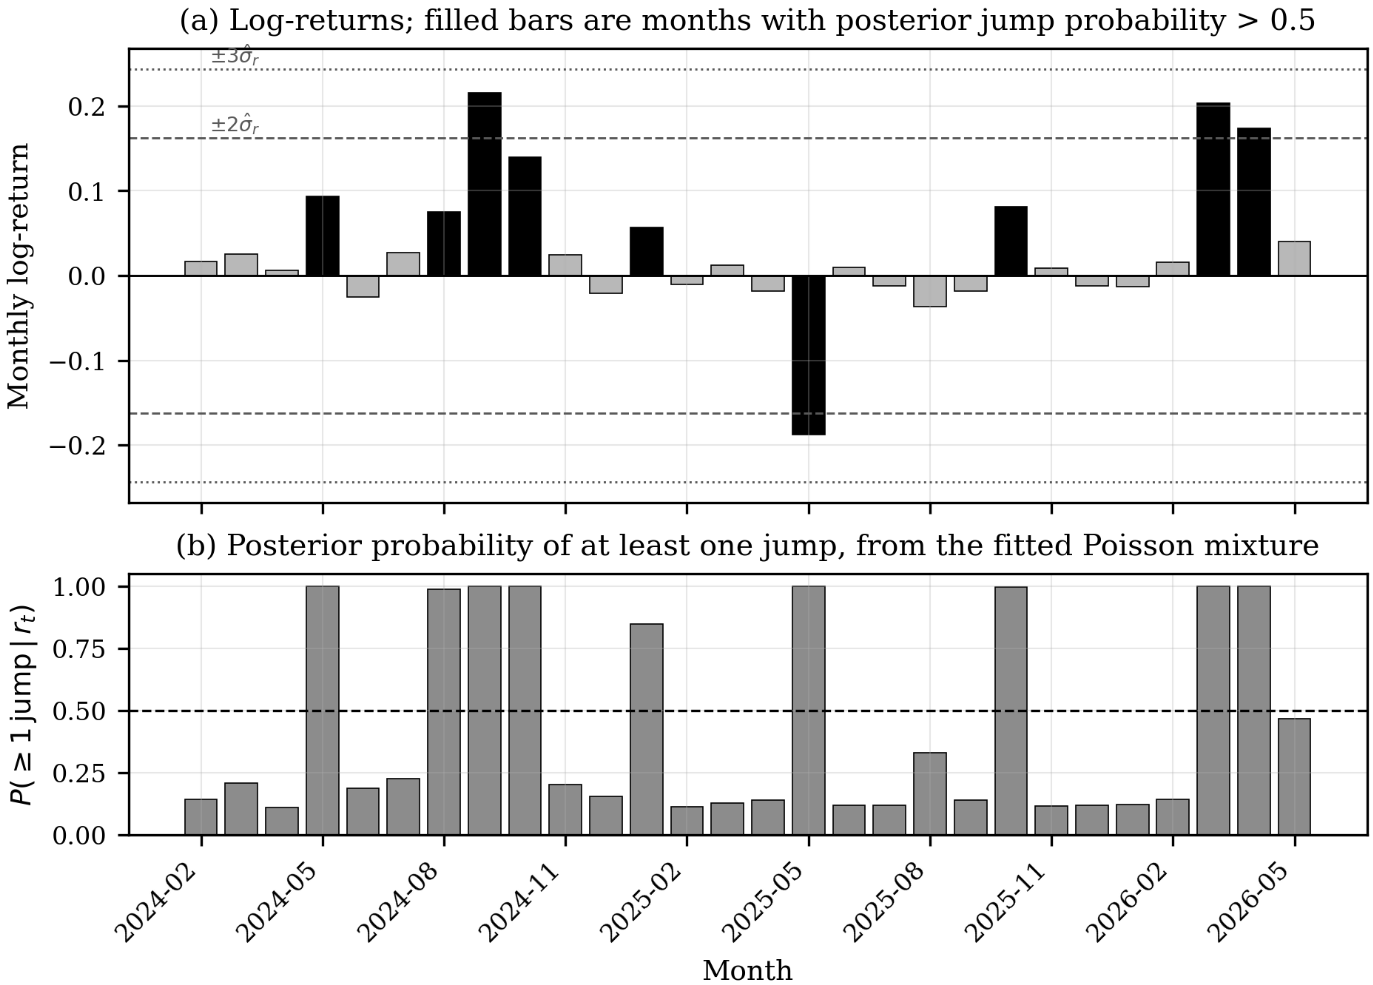

Figure 4: (a) Monthly log-returns of the PMS series with two- and three-standard-deviation reference bands; filled bars denote months whose posterior jump probability exceeds one half. (b) Posterior probability that each month contains at least one jump, computed from the fitted Poisson mixture using equation (9). The separation between ordinary and jump months is sharp.

dependence structure, and the surrogate is therefore not a generally identified operator of the price process; its accuracy claims are restricted to the jump-diffusion scenario family and the sampling procedure of Section 3.2.4 on which it is trained and tested, and it is used only as a fast interpolant within that family, not as a substitute for the exact solver outside it. Training data are generated by random sampling of feasible capacity pairs, each evaluated on an independent subsample of 400 scenarios; unlike the drift problem, here training data are unlimited in principle because the exact solver can be queried at will. The network has two hidden layers of twenty-four units and is trained by Adam for 6000 iterations on a 75/25 train-test split.

4. Results

4.1. Descriptive evidence

Figure 3 shows the two published price series. The PMS series rose from 668.30 naira per litre in January 2024 to a local peak of 1261.65 naira in March 2025, fell to 970.59 naira by September 2025 during a period of intensified downstream competition, then rose again to 1596.25 naira by May 2026. The AGO series averaged 1457.41 naira per litre with a standard deviation of 159.18 naira over the seventeen verified months. Both series display flat stretches punctuated by steps. Figure 4(a) plots the monthly log-returns against two- and three-standard-deviation bands; several returns lie far outside the bands and are not evenly distributed through the sample but cluster around identifiable events.

4.2. Model selection

Table 3 reports the estimation results. The ranking is consistent across all three criteria: the Merton jump-diffusion attains the lowest Akaike criterion at −67.12, the lowest corrected criterion at −64.39 and the lowest Bayesian criterion at −60.46. The

Table 3: Maximum-likelihood estimation and model selection for the four candidate price processes.

Model	Parameters k	Log-likelihood	AIC	AIC$_c$	BIC	Rank
M1 GBM	2	31.112	-58.22	-57.74	-55.56	3
M2 OU	3	31.412	-56.82	-55.82	-52.83	4
M3 MJD	5	38.558	-67.12	-64.39	-60.46	1
M4 MRJD	6	39.333	-66.67	-62.67	-58.67	2

M1 geometric Brownian motion; M2 Ornstein-Uhlenbeck; M3 Merton jump-diffusion; M4 mean-reverting jump-diffusion. Rank is by AIC$_c$; lower is better.

Table 4: Model comparison. The two jump comparisons are regular nested tests whose significance is assessed by a boundary-aware parametric bootstrap; the two mean-reversion comparisons are non-nested and are assessed by the difference in AIC.

Comparison	Added feature	Nested?	LR statistic	Bootstrap p	ΔAIC	Conclusion
M1 GBM vs M3 MJD	Jumps	Yes	14.89	0.025	-	Supported
M2 OU vs M4 MRJD	Jumps	Yes	15.84	0.040	-	Supported
M1 GBM vs M2 OU	Mean reversion	No	-	-	-1.40	Not supported
M3 MJD vs M4 MRJD	Mean reversion	No	-	-	-0.45	Not supported

Each bootstrap null distribution is simulated from the corresponding fitted no-jump model: the M1-vs-M3 test uses 199 geometric-Brownian-motion replicates (95th percentile 9.66) and the M2-vs-M4 test uses 99 Ornstein-Uhlenbeck replicates (95th percentile 11.05). A negative ΔAIC favours the model without mean reversion.

Table 5: Parameter estimates for the selected Merton jump-diffusion (M3).

Name	Symbol	Estimate	Standard error	z-statistic	Interpretation
Drift	μ	0.0141	0.0757	0.19	Diffusive drift (per year)
Volatility	σ	0.0706	0.0168	4.21	Diffusive volatility (per $\sqrt{\text{year}}$)
Intensity	λ	6.6538	2.7908	2.38	Jump intensity (jumps per year)
Jump mean	μ_J	0.0540	0.0304	1.77	Mean jump size (log price)
Jump dispersion	σ_J	0.0962	0.0246	3.91	Jump size standard deviation
Log-price drift	$\mu + \lambda\mu_J$	0.3731	-	-	Annualised drift of $\log S_t$
Expected jumps/month	$\lambda\Delta$	0.5545	-	-	Poisson mean over one month
P(at least one jump)	$1 - e^{-\lambda\Delta}$	0.4256	-	-	Probability of ≥ 1 jump in a month

geometric Brownian motion is second on the Bayesian criterion but well behind on the others, the Ornstein-Uhlenbeck process is worst on every criterion, and the full mean-reverting jump-diffusion attains the highest log-likelihood, as it must since it has the most parameters, but its extra parameter is not repaid: its corrected criterion of −62.67 is worse than that of the simpler jump model.

The likelihood-ratio tests in Table 4 sharpen the point for the two valid nested comparisons. Adding jumps to the geometric Brownian motion yields a statistic of 14.89, and adding jumps to the Ornstein-Uhlenbeck process yields 15.84. Because $\lambda = 0$ is a boundary point, we assess significance by parametric bootstrap rather than by the χ^2 distribution, simulating from each fitted no-jump model and refitting both specifications on every replicate. For the geometric Brownian motion null the observed statistic of 14.89 gives a bootstrap p-value of 0.025 ($B = 199$), and for the Ornstein-Uhlenbeck null the observed statistic of 15.84 gives a bootstrap p-value of 0.04 ($B = 99$). The jump term is therefore supported in both comparisons even under the correct boundary-aware reference distribution. For mean reversion, which is not a regular nested restriction here, the information criteria are decisive in the other direction: the Akaike criterion of the geometric Brownian motion is 1.40 units below that of the Ornstein-Uhlenbeck process, and that of the jump-diffusion is 0.45 units below the mean-reverting jump-diffusion, so adding mean reversion does not improve the fit.

Figure 5 provides distributional confirmation: in panel (a) the fitted Gaussian density is visibly too narrow in the centre and too thin in the tails while the Poisson mixture tracks both features, and in panel (b) the quantile-quantile plot shows the Gaussian comparison bending away from the diagonal at both extremes while the mixture remains close to it. Figure 6 displays the information criteria side by side.

4.3. Jump identification and external validation

Table 6 lists the nine months whose posterior jump probability exceeds one half, together with the market events recorded for those months, and the agreement is close. The largest positive move, a log-return of +0.216 (a 24.1 per cent price rise) in September

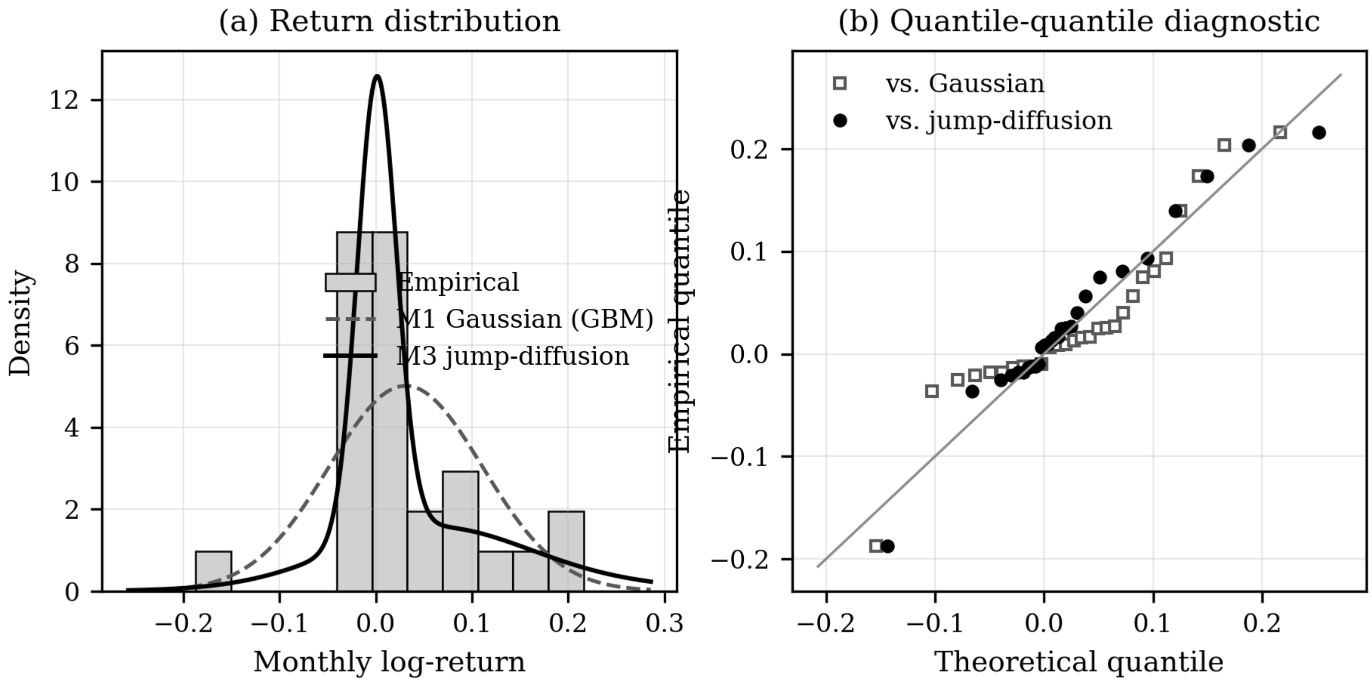

Figure 5: Distributional diagnostics for the monthly log-returns. (a) Empirical histogram against the fitted Gaussian density of model M1 and the Poisson-mixture density of model M3. (b) Quantile-quantile comparison of the empirical returns against both fitted distributions, with the forty-five degree reference line.

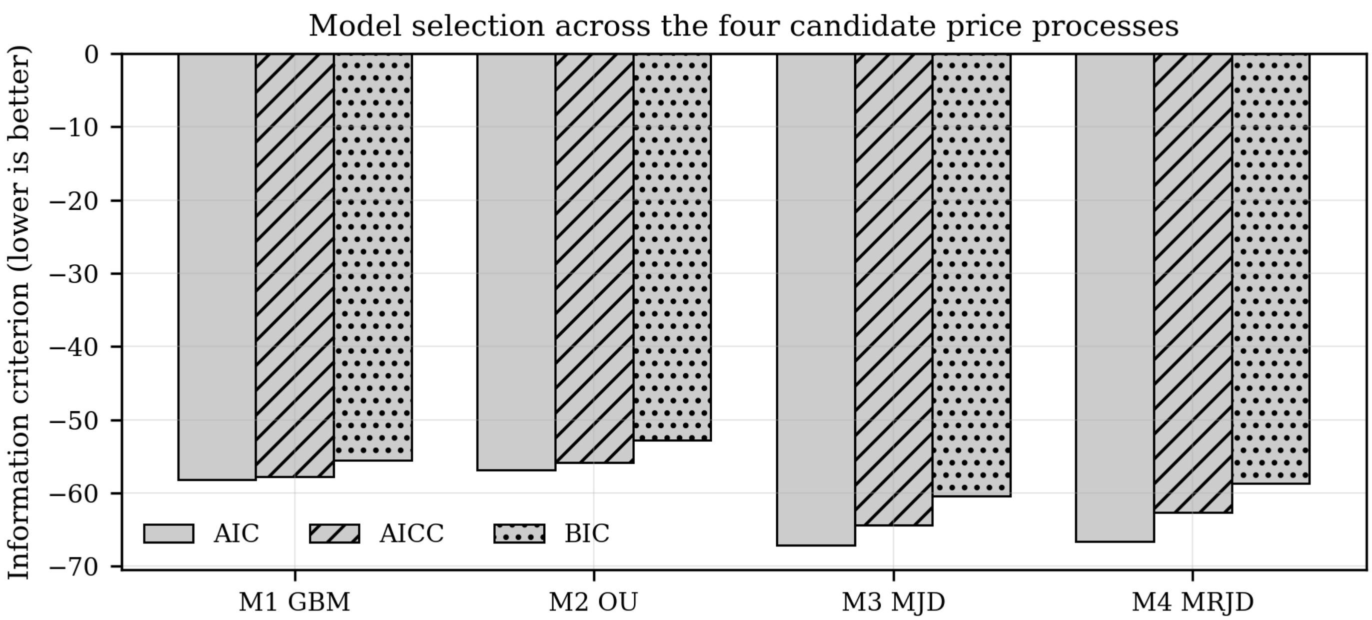

Figure 6: Akaike, corrected Akaike and Bayesian information criteria for the four candidate price processes. Lower values indicate better models; the Merton jump-diffusion M3 is preferred on all three criteria.

2024 with posterior probability 1.00, coincides with the entry of domestically refined petrol and the associated repricing, when the national oil company moved its pump price to about 897 naira per litre; it is followed by a +0.140 log-return (a 15.0 per cent rise) in October 2024. The single largest fall, a log-return of −0.187 (a 17.1 per cent price decline) in May 2025 with probability 1.00, coincides with the period of intense downstream competition that pulled the national average from 1239.33 naira in April to 1027.76 naira in May. The +0.081 log-return (an 8.4 per cent rise) in October 2025 with probability 1.00 coincides with the approval of a 15

Table 6: Months identified as containing a jump, with posterior probability and the recorded market event.

Month	Log-return	Price change (%)	$P(\geq 1 \text{ jump} \mid r_t)$	Recorded market event
2024-05	+0.0930	+9.75	1.000	Post-deregulation price adjustment
2024-08	+0.0749	+7.78	0.990	Continued upward repricing
2024-09	+0.2158	+24.08	1.000	Entry of domestically refined petrol; national price near 897 NGN/L
2024-10	+0.1396	+14.98	1.000	Continuation of the repricing sequence
2025-01	+0.0566	+5.82	0.849	New-year price adjustment
2025-05	-0.1872	-17.07	1.000	Intensified downstream price competition
2025-10	+0.0808	+8.42	0.996	15 per cent ad-valorem import duty on PMS and AGO
2026-03	+0.2033	+22.55	1.000	Renewed escalation in pump prices
2026-04	+0.1737	+18.97	1.000	Continued escalation

Table 7: Optimal investment, cost and risk as a function of the investment budget ($\beta = 1$, $\alpha = 0.95$).

Budget B (m)	y_{pv} (kWp)	y_b (kWh)	Installed (m)	$\mathbb{E}[Q]$ (m)	$\text{CVaR}_{0.95}$ (m)	Total (m)	Saving (%)	Payback (yr)	Shadow
0 (status quo)	0.00	0.0	0.00	3.818	5.582	3.818	-	-	-
1	1.00	0.0	0.90	2.604	3.853	2.874	24.7	0.74	2.67
2	2.00	0.0	1.80	1.608	2.401	2.148	43.7	0.81	2.18
3	3.00	0.5	2.92	0.927	1.380	1.809	52.6	1.01	1.36
4	3.25	2.0	3.83	0.760	1.132	1.926	49.6	1.25	0.13
5	3.50	3.5	4.72	0.625	0.929	2.075	45.7	1.48	0.05
6	3.75	4.5	5.40	0.523	0.774	2.184	42.8	1.64	0.05
7	4.25	7.0	6.97	0.322	0.468	2.479	35.1	2.00	0.01
8	4.25	7.0	6.97	0.322	0.468	2.479	35.1	2.00	-0.00

per cent ad-valorem import duty on petrol and diesel, and the two large moves at the end of the sample, log-returns of +0.203 and +0.174 (price rises of 22.6 and 19.0 per cent) in March and April 2026, correspond to the renewed escalation that carried the price from 1051.47 naira in February to 1532.93 naira in April.

That a purely statistical filter, given no event information, recovers a set of months matching the documented policy record is meaningful external validation: it supports interpreting the estimated jump component as the signature of discrete policy and market-structure events rather than an artefact of a flexible functional form. Figure 4(b) plots the posterior probabilities, and the separation is sharp, with most months below 0.25 and the identified jump months above 0.85.

4.4. Scenario analysis

Figure 7 presents the scenario fan for the twelve months following May 2026. Starting from 1596.25 naira per litre, the median simulated price reaches 2270.6 naira after twelve months, with a fifth percentile of 1506.6 naira and a ninety-fifth percentile of 3820.2 naira. The distribution is strongly right-skewed, since the mean of 2417.0 naira exceeds the median by 6.4 per cent, a direct consequence of the positive mean jump size and exactly the feature a symmetric model would miss. The ninety-fifth percentile is about 2.5 times the fifth percentile, so a firm planning on the median outcome is planning on a scenario it has only even odds of experiencing while the adverse tail lies far above it. The figure also shows a zero-drift stress case, constructed by removing the drift while retaining the jump structure, in which the median twelve-month price is 1664.3 naira; even with no drift the distribution remains wide, because jump risk alone generates substantial dispersion.

4.5. Optimization results

Status quo. With no investment the representative enterprise faces an expected annual energy cost of 3,818,401 naira, equivalent to 530.3 naira per kilowatt-hour. Of its annual demand, 50.7 per cent is met from the grid, 48.1 per cent from the generator and 1.2 per cent goes unserved. The conditional value-at-risk of operating cost at the 95 per cent level is 5,581,569 naira, that is 46.2 per cent above the expected value. The unit cost of 530.3 naira per kilowatt-hour against a Band A tariff of 209.50 naira per kilowatt-hour quantifies the penalty of self-generation: the firm pays about 2.5 times the regulated tariff for its blended energy.

Optimal investment. Table 7 reports the budget sweep. At the risk-aversion weight $\beta = 1$ and a budget of six million naira, the optimum is 3.75 kilowatts-peak of photovoltaic capacity with 4.5 kilowatt-hours of storage, an installed cost of 5,400,000 naira. Expected operating cost falls to 522,266 naira and the conditional value-at-risk of operating cost to 774,000 naira. Including the annualised capital charge of 1,661,552 naira, the total expected annual cost is 2,184,421 naira, a reduction of 42.8 per cent against the status quo. Measured on operating cost, the quantity that carries the price risk, the conditional value-at-risk falls by 86.1 per cent; this figure refers to operating cost only and excludes the deterministic capital charge, which is fixed and carries no risk. The risk reduction exceeds the expected-cost reduction because solar capacity displaces the one input whose price is stochastic, converting a random cost into a largely deterministic tariff-and-capital cost.

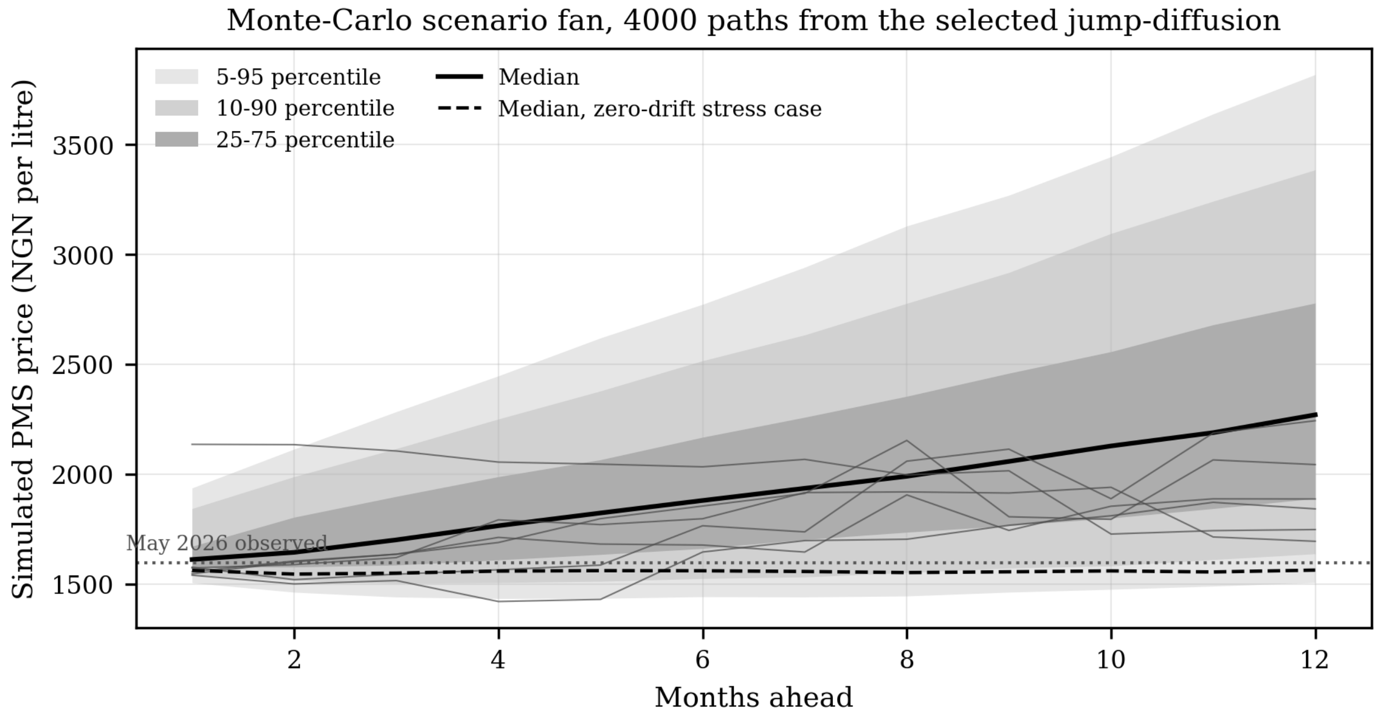

Figure 7: Monte-Carlo scenario fan for the twelve months following May 2026, generated from the selected jump-diffusion using equation (10) with 4000 paths. Shaded bands give the 5-95, 10-90 and 25-75 percentile ranges; thin lines show six individual sample paths. The dashed line is the median of a zero-drift stress case retaining the jump structure.

All monetary values in million naira. Total is the annualised capital charge plus expected operating cost; CVaR is of operating cost. The shadow price is the reduction in the mean-CVaR objective per additional naira of budget. The cost-minimising budget is 3 million naira.

Capital rationing within the model. Table 7 also reports the shadow price of investment capital, computed as the reduction in the mean-CVaR objective per additional naira of budget. At a budget of one million naira it is 2.67, meaning each additional naira of budget returns 2.67 naira of annual objective improvement; it falls to 2.18 at two million and 1.36 at three million, then to 0.13 at four million and effectively to zero beyond seven million. The lowest-total-cost solution is not the largest feasible one: it occurs at a budget of three million naira, with 3.0 kilowatts-peak and 0.5 kilowatt-hours, delivering a 52.6 per cent reduction in expected total cost with a simple payback of 1.01 years, after which additional capital raises the annualised charge faster than it lowers operating cost. Within the model, therefore, the budget constraint (14) rather than the objective (13) limits the attainable saving. We phrase this as a statement about the model: a positive shadow price shows the value of relaxing the assumed budget, and identifying capital as the real-world binding constraint for Nigerian MSMEs would require firm-level financing data that this study does not use. The direction is nonetheless consistent with the survey finding of Achmad and Subriadi that inadequate capital obstructs MSME technology adoption [22] and with Abba and colleagues on risk barriers to decentralised renewable investment in Nigeria [26]. Figure 8 plots the cost-risk frontier traced by the budget, and Figure 9 shows the objective surface and the capacity response with the shadow price.

Storage as a risk instrument. Table 8 reports the frontier traced by β at a fixed six-million-naira budget. At $\beta = 0$ the optimum is 2.75 kilowatts-peak with zero storage; storage enters only at $\beta = 0.25$ and grows to 5.0 kilowatt-hours at $\beta = 1.5$, while photovoltaic capacity rises more gently from 2.75 to 4.0 kilowatts-peak. Storage is therefore not economic on expected-cost grounds at current Nigerian prices; it is purchased to compress the upper tail of the cost distribution. Between $\beta = 0$ and $\beta = 1$ the conditional value-at-risk of operating cost falls from 1,553,000 to 774,000 naira, a reduction of 50.2 per cent. A risk-neutral appraisal of a battery subsidy would reject it; a risk-aware appraisal of the same subsidy, recognising that a small firm's survival depends on its worst year rather than its average year, would not.

Energy mix. At the $\beta = 1$ optimum the annual energy mix shifts from 50.7 per cent grid, 48.1 per cent generator and 1.2 per cent unserved to 72.2 per cent solar, 25.5 per cent grid, 2.3 per cent generator and zero unserved. Generator dependence falls by a factor of about twenty, and the elimination of unserved energy carries a productivity benefit that the cost figures alone understate.

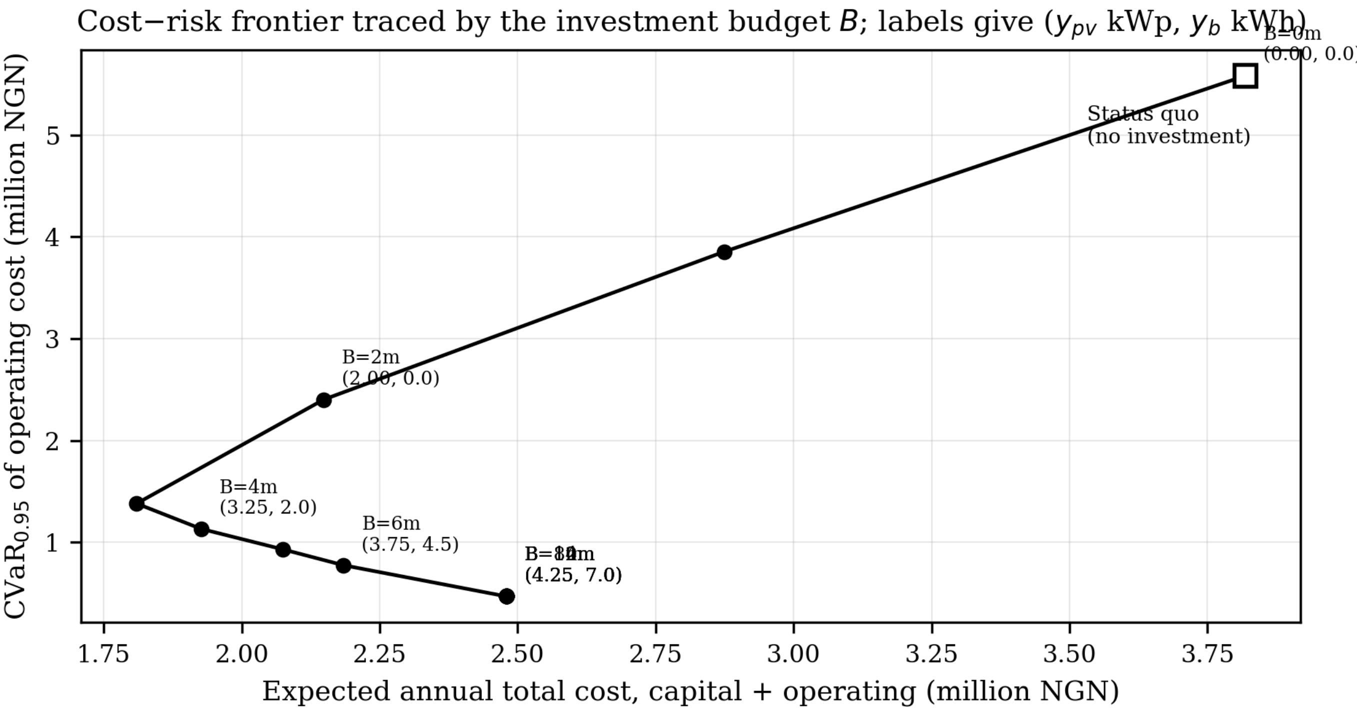

Figure 8: Cost-risk frontier traced by the investment budget at risk-aversion weight one. Each point is the optimal solution at that budget, labelled with the optimal photovoltaic and storage capacities. The open square marks the status quo with no investment.

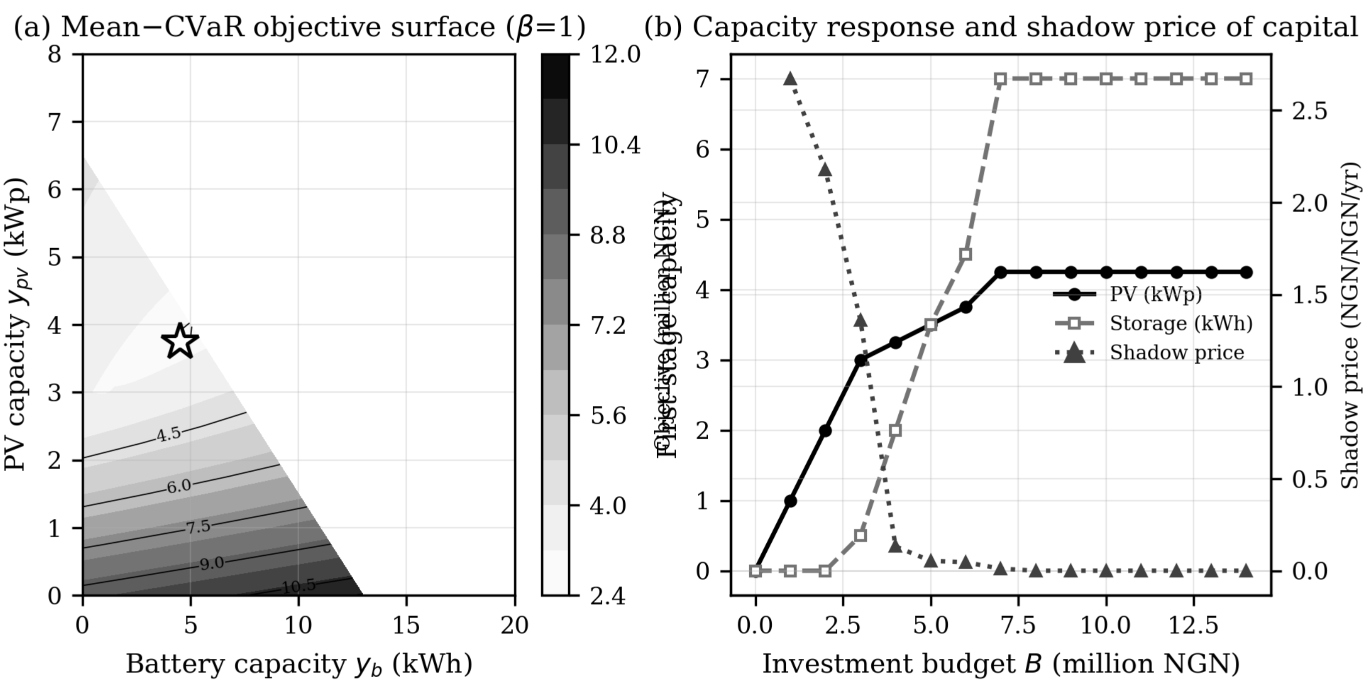

Figure 9: (a) Mean-CVaR objective surface over the feasible capacity lattice at a six-million-naira budget; the star marks the optimum and the blank upper-right region is excluded by the budget constraint (14). Contour lines are labelled in million naira. (b) Optimal first-stage capacities and the shadow price of investment capital as functions of the budget.

Table 8: Cost-risk frontier traced by the risk-aversion weight β at a fixed budget of 6 million naira.

β	y_{pv} (kWp)	y_b (kWh)	Installed (m)	$\mathbb{E}[Q]$ (m)	$\mathrm{VaR}_{0.95}$ (m)	$\mathrm{CVaR}_{0.95}$ (m)
0.00	2.75	0.0	2.48	1.043	1.399	1.553
0.25	3.00	0.5	2.92	0.927	1.249	1.380
0.50	3.25	2.0	3.83	0.760	1.019	1.132
1.00	3.75	4.5	5.40	0.523	0.701	0.774
1.50	4.00	5.0	5.85	0.480	0.642	0.709
2.00	4.00	5.0	5.85	0.480	0.642	0.709

Storage is absent from the risk-neutral optimum and enters only as risk aversion increases.

Table 9: In-sample and out-of-sample performance of the three drift models.

Model	Type	Training RMSE	Hold-out RMSE	Hold-out price MAPE (%)	Verdict
Merton jump-diffusion	Parametric	0.0749	0.0950	6.75	Best generalisation
Universal differential equation	Grey-box	0.0541	0.2283	16.58	Intermediate
Multilayer perceptron	Black-box	0.0034	0.2451	19.58	Severe overfitting

Models are trained on the first 22 transitions and evaluated on the final 6. RMSE is in units of the monthly log-return.

Table 10: Deep recourse operator: accuracy and computational performance.

Metric	Expected cost $\mathbb{E}[Q]$	$\mathrm{CVaR}_{0.95}[Q]$
Coefficient of determination R^2	0.9996	0.9968
Mean absolute percentage error (%)	0.92	2.47
Design points generated	1200	1200
Hold-out points	300	300
Exact evaluation time (ms)	1.194	1.194
Surrogate evaluation time (ms)	0.0119	0.0119
Speed-up factor	100×	100×

4.6. Scientific machine learning results

Drift learning: a negative result. Table 9 reports the comparison. On the twenty-two training transitions the black-box network attains a root-mean-square error of 0.0034, more than twenty times better than the parametric model's 0.0749 and better than the grey-box model's 0.0541. On the six held-out transitions the ordering reverses: the parametric jump-diffusion attains 0.0950, the grey-box model 0.2283, and the black-box network the worst value of 0.2451. Translated into price levels, the out-of-sample mean absolute percentage errors are 6.75 per cent for the parametric model, 16.58 per cent for the grey-box model and 19.58 per cent for the black-box network. This is textbook overfitting, reported because it is the honest result: the black-box network reduced its training error by a factor of twenty-two while increasing its generalisation error. The grey-box formulation, which anchors the drift to the selected model and learns only a residual, mitigates the failure but does not eliminate it, because with twenty-eight transitions there is not enough information to identify a residual drift function on top of an already adequate parametric structure. Figure 10(a) shows the training curves and Figure 11 the resulting comparison. Where a national statistical series is monthly and a structural break is recent, the number of usable observations is measured in tens, and flexible approximation of the data-generating process is actively harmful; in this application the correct decision was to let the mechanistic model do the work.

Recourse surrogate: a positive result. The picture changes entirely for the second task. Table 10 reports the accuracy and computational performance of the surrogate. With the expanded feature set of equation (25) the deep recourse operator attains a coefficient of determination of 0.9996 for expected cost and 0.9968 for conditional value-at-risk on 300 held-out design points, with mean absolute percentage errors of 0.92 and 2.47 per cent respectively; Figure 12(a) plots predicted against exact values, which lie tightly on the diagonal. An exact evaluation requires 1.19 milliseconds, dispatching 4000 scenarios over twelve months, while the surrogate returns both risk measures in 0.012 milliseconds, a speed-up of about 100 times, shown in Figure 12(b). The contrast between the two applications is the paper's methodological lesson: drift learning failed because training data were scarce and fixed by history, whereas surrogate learning succeeded because training data could be generated on demand from an exact solver. The right question before applying machine learning is not whether the problem is complex but whether the training data are abundant.

4.7. Sensitivity analysis

Table 11 and Figure 13 report the sensitivity of the optimal solution to nine parameters, each varied to a low and a high value around its base case. To avoid the ambiguity of measuring an expected-cost saving while the optimiser trades expected cost against risk, the saving here is measured on the mean-CVaR objective that is actually minimised; because that objective is non-increasing in

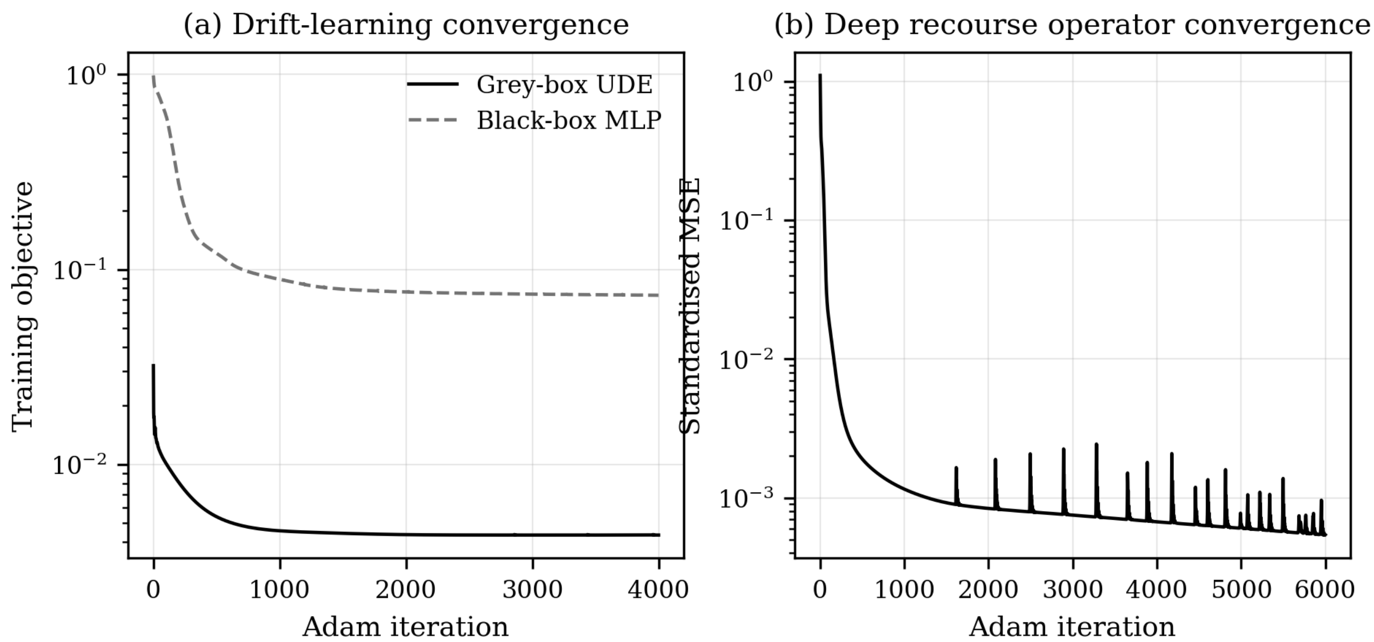

Figure 10: Convergence of the two scientific machine learning components. (a) Training objective of the grey-box universal differential equation, whose mechanistic backbone is the selected-model drift, and of the black-box network for drift learning. (b) Standardised mean squared error of the deep recourse operator.

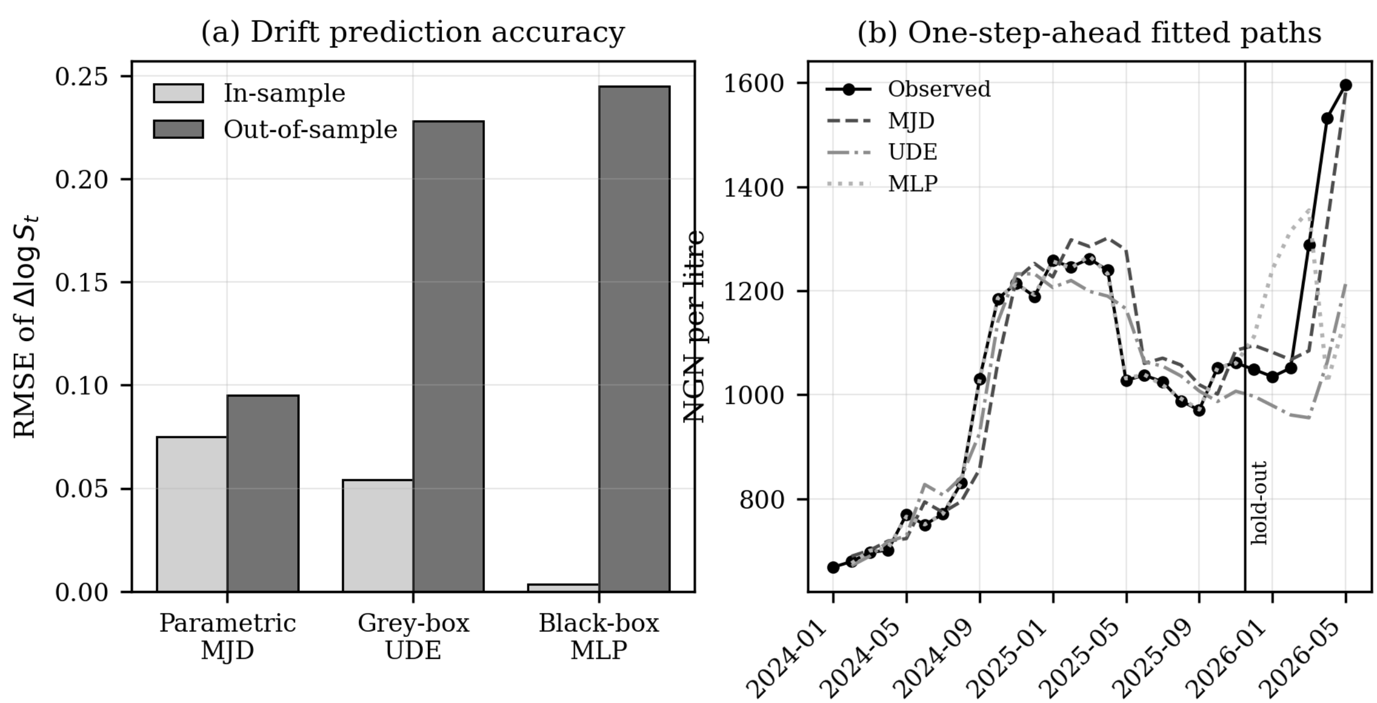

Figure 11: Comparison of the three drift models. (a) Root-mean-square error of the predicted monthly log-return, in sample and on the six withheld months. (b) One-step-ahead fitted price paths against the observed series; the vertical line marks the start of the hold-out period. The black-box network attains the best training fit and the worst generalisation.

every cost parameter, the saving now moves in the intuitive direction throughout, and it exceeds the expected-cost saving of Table 7 because it also credits the reduction in tail risk. The base-case objective saving is 68.2 per cent, and across all eighteen perturbations it remains between 63.9 and 72.6 per cent, so no parameter setting reverses the finding that investment is worthwhile. The optimal

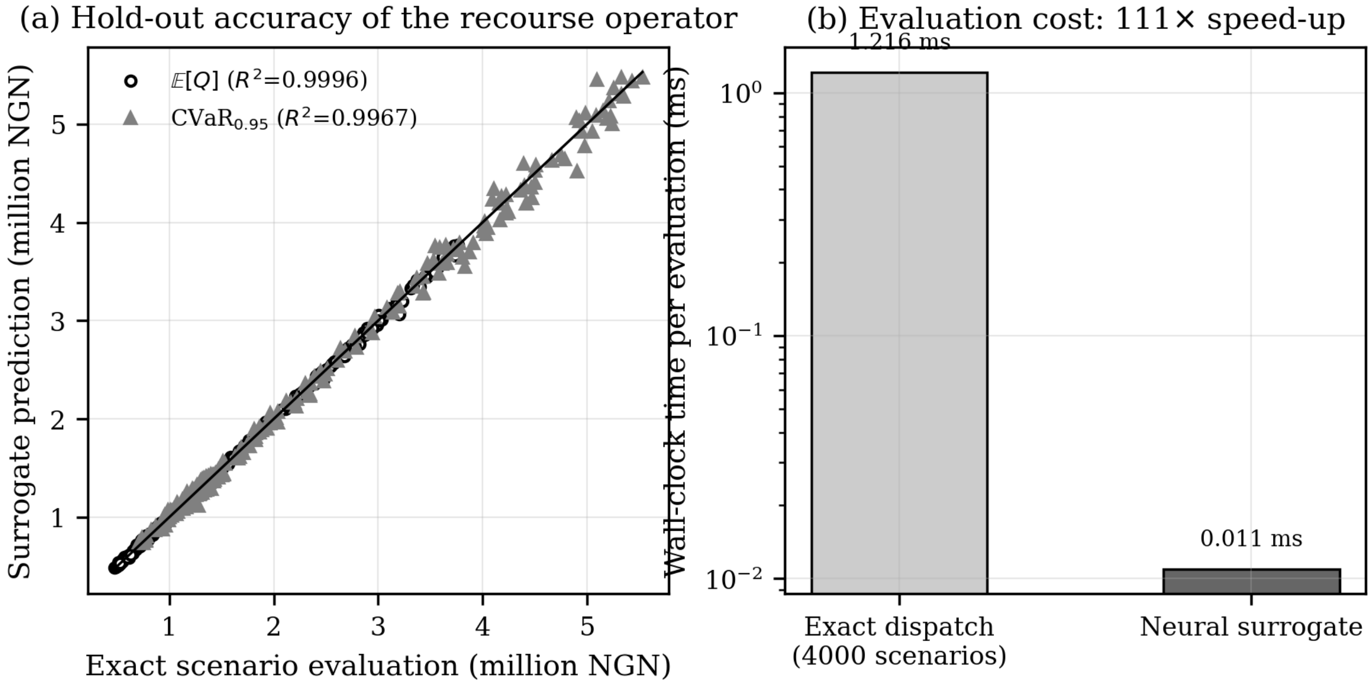

Figure 12: Deep recourse operator performance. (a) Surrogate predictions against exact scenario evaluations for expected cost and conditional value-at-risk on held-out design points. (b) Wall-clock evaluation time on a logarithmic scale, showing the speed-up over exact dispatch of 4000 scenarios.

Table 11: Sensitivity of the optimal mean-CVaR objective saving to model parameters.

Parameter	Low	Base	High	Saving at low (%)	Saving at base (%)	Saving at high (%)	Range (pp)
Cost of capital	0.18	0.25	0.32	71.7	68.2	65.0	6.7
Genset efficiency (L/kWh)	0.34	0.42	0.5	63.9	68.2	71.5	7.6
Battery cost (NGN/kWh)	350,000	450,000	550,000	70.7	68.2	66.9	3.8
PV capital cost (NGN/kWp)	700,000	900,000	1,100,000	71.2	68.2	65.7	5.5
Peak sun hours (h/day)	4.4	4.9	5.4	67.0	68.2	69.6	2.6
Daytime coincidence share	0.45	0.55	0.65	64.2	68.2	72.6	8.4
Value of lost load (NGN/kWh)	1,000	1,500	2,200	67.8	68.2	68.7	0.9
Band A tariff (NGN/kWh)	180	209.5	240	68.5	68.2	67.9	0.6
Mean grid availability	0.4	0.55	0.7	68.2	68.2	68.2	0.0

capacities and objective values underlying every case are reported in Appendix A (Table **??**), so that the re-optimisation can be checked directly.

Saving is measured on the mean-CVaR objective that is minimised, so it is monotone in every cost parameter and exceeds the expected-cost saving of Table 7.

The cost of capital is the most influential parameter: reducing it from 25 to 18 per cent raises the saving from 68.2 to 71.7 per cent, while raising it to 32 per cent lowers the saving to 65.0 per cent, because the discount rate enters the annualised charge on the same physical asset. This reinforces the capital-rationing result from a second direction. Generator efficiency is the next most influential: a less efficient generator, consuming 0.50 rather than 0.42 litres per kilowatt-hour, raises the saving to 71.5 per cent because the displaced alternative becomes more expensive. The daytime coincidence share φ has a comparable effect: moving it from 0.45 to 0.65 raises the saving from 64.2 to 72.6 per cent, since a larger share of demand coincident with sunlight makes photovoltaic capacity more valuable. Because the assumed base value of 0.55 is not derived from metered firm-level data, this sensitivity is reported explicitly: the qualitative conclusion is preserved across the plausible range, while the magnitude of the gain is moderately sensitive to the coincidence share. Photovoltaic capital cost and peak sun hours have moderate effects. The electricity tariff has remarkably little effect: varying the Band A rate from 180 to 240 naira per kilowatt-hour changes the saving only from 68.5 to 67.9 per cent and does not change the optimal capacities, and the same is true of mean grid availability. Because the generator, not the grid, is the marginal supplier, and because the generator's cost is set by a jump-driven fuel price, moderate tariff adjustments barely register: interventions aimed at the regulated tariff will not materially change the economics of enterprise-level energy resilience, whereas interventions aimed at the cost of capital will.

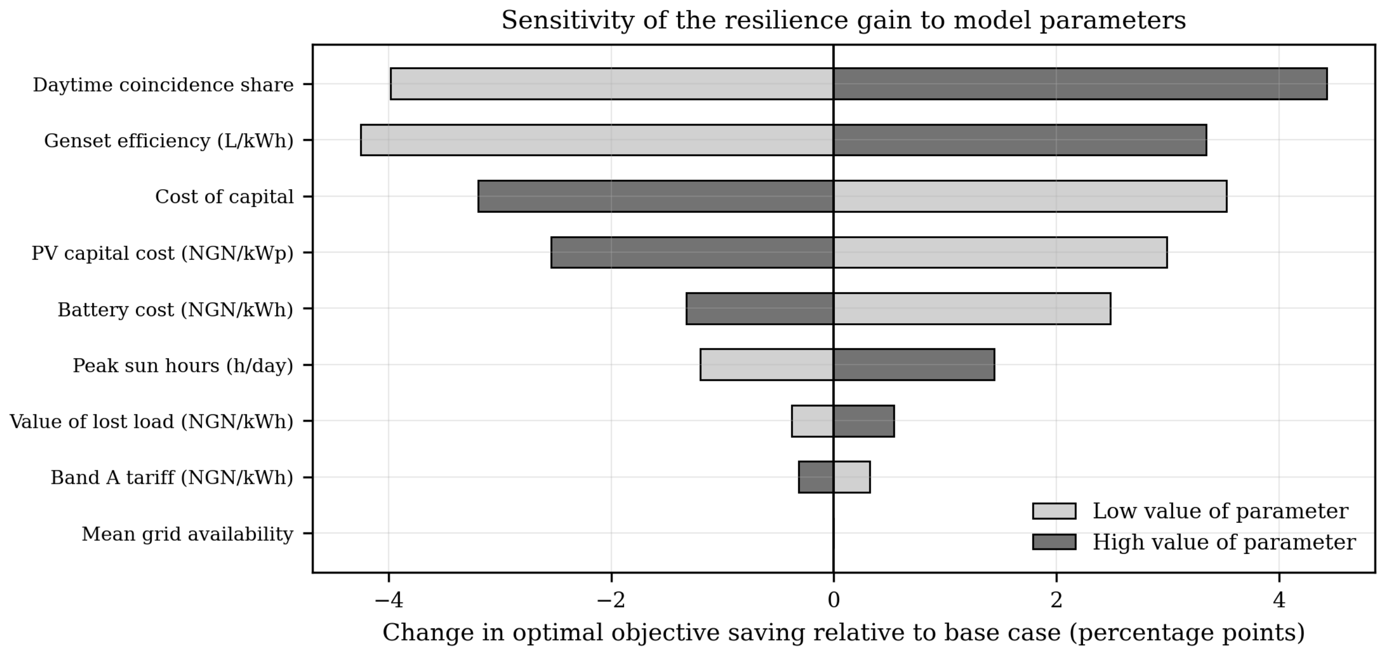

Figure 13: Tornado diagram of the sensitivity of the optimal mean-CVaR objective saving to nine model parameters, each perturbed to a low and a high value about its base case. Bars show the change in percentage points relative to the base case. The cost of capital is the most influential parameter; the electricity tariff is among the least influential.

5. Discussion

5.1. Interpretation

Three findings carry the argument. The first is that, for this sample, the Nigerian post-subsidy pump price is best described as jump-dominated with no detectable mean reversion. Under a mean-reverting model a firm experiencing a cost shock should absorb it and wait for reversion, because the expected future price is anchored; our estimates provide no evidence of such an anchor, and the entire drift comes from jumps whose mean size is positive. A firm that waits is not waiting out a deviation but absorbing a largely permanent increase, so the optimal response is structural rather than financial, which is what the optimization results confirm.

The second is that, within the model, capital access rather than technology cost or tariff policy is the binding constraint. The shadow price of about 2.67 naira per naira per year at low budgets, together with a payback of about one year at the cost-minimising scale, describes an investment that would be made immediately by any firm able to finance it. We are careful not to over-claim here: the model imposes a budget but does not observe real firms' financing limits, so the result establishes the value of relaxing the assumed budget rather than proving that finance is the operative real-world constraint. The sensitivity analysis independently identifies the cost of capital as the most influential parameter, which points in the same direction.

The third is that the two roles of scientific machine learning must be distinguished. Our results show the same family of methods failing at one task and succeeding at another within a single study, and the discriminating variable is data abundance rather than problem difficulty.

5.2. Policy implications

The findings support four propositions, stated in decreasing order of confidence. First, concessional finance for enterprise-scale solar is likely to be more effective per naira than either fuel or tariff subsidy, since a three-million-naira facility to a single representative firm removes about two million naira of annual energy cost and pays back in roughly a year. Second, storage subsidies should be justified on risk grounds, not average-cost grounds, because storage is not economic at $\beta = 0$ but becomes so as risk aversion rises, so an expected-value appraisal will systematically reject storage support that a risk-aware appraisal would accept. Third, tariff instruments are weak levers for this population, because the generator is the marginal supplier for most firms and the regulated tariff is largely inframarginal. Fourth, price stability has measurable value independent of price level, since the zero-drift stress scenario retains wide dispersion from jump risk alone, so measures that reduce the frequency or size of discrete price events would deliver benefits even without lowering the average price. These propositions follow from the model under its stated assumptions and for the representative firm of Table 2; they are quantitative illustrations of a mechanism rather than costed recommendations.

5.3. Limitations

We state the limitations plainly, since several are material.

Sample size. Twenty-nine monthly observations is a short series. It is the complete published record for the post-subsidy regime at national monthly frequency, and a longer window would splice together two structurally different policy regimes, but it remains short. The estimated jump intensity of 6.65 per year has a standard error of 2.79, so a plausible range runs from roughly one to twelve jumps per year, and the mean jump size is only marginally significant. Our conclusions rest on the qualitative structure of the process rather than on precise coefficient values, and the failure to detect mean reversion should be read as an absence of evidence at this sample size rather than as proof of its absence. Because the downstream optimization is driven by the fitted process, this parameter uncertainty propagates into the scenario fan and, through it, into the optimal capacities; the sensitivity analysis of Section 4.7 is the appropriate guard, and it shows the qualitative investment conclusion to be robust across the tested ranges.

Representative-firm design. We model one enterprise archetype. Nigerian MSMEs are heterogeneous in load, sector, location, grid band and financial capacity, and the optimum will differ across them; the reported figures of a 42.8 per cent cost reduction, an 86.1 per cent operating-cost tail-risk reduction and a 2.67 naira shadow price are properties of this representative firm and should not be read as population averages or obtained by multiplying by the number of firms. Delineating the scope precisely: the framework applies to a self-generating firm whose marginal energy source is a fuel generator and whose daytime load share is near the assumed value; firms that are predominantly grid-served, or whose load is concentrated at night, would show smaller gains.

No firm-level data. We use published national statistics rather than surveyed enterprise data. This is a deliberate design choice that makes the framework applicable where no firm-level data system exists, but it means the demand profile, the daytime coincidence share φ and the value of lost load are assumptions rather than measurements. The daytime coincidence share of 0.55 is particularly influential, as Section 4.7 shows, and would benefit from metered validation; we retain 0.55 as a central value for a firm with a normal working day and report the full sensitivity rather than relying on the point value.

Independence assumptions. Fuel price and grid availability are modelled as independent. In reality gas-supply constraints and foreign-exchange pressure can affect both, so the true dependence is plausibly positive; a positive dependence would make adverse months worse than we model, since expensive-fuel months would coincide with poor grid supply, and would therefore strengthen rather than weaken the case for on-site solar investment. Jump sizes are likewise assumed independent and identically distributed, whereas policy events plausibly cluster.

Energy-balance recourse. The recourse of Section 3.2.6 is a monthly energy-balance model. It does not track a battery state of charge, charge and discharge power limits, chronology within the month, or carry-over between months, and it treats monthly grid-available energy and battery throughput as fully allocatable to residual demand. These simplifications can overstate the service that a given capacity delivers, so the reported savings should be read as an upper envelope for the modelled configuration; an hourly formulation with explicit state-of-charge dynamics would refine, and probably moderate, the storage valuation.

Deterministic technical performance and partial equilibrium. Photovoltaic yield, battery degradation and equipment failure are treated deterministically, and degradation in particular would reduce realised savings over an asset life; reliability modelling of the kind developed by Shen and colleagues [33] and Kumar and co-authors [38] would refine the appraisal. Finally, we take prices as exogenous to the firm, so the framework cannot capture any effect of large-scale solar adoption on fuel demand or prices.

5.4. Comparison with existing work

Our optimization architecture is closest in spirit to Moghadaspoor and colleagues [10] and Mozaffari and colleagues [11], both of which use two-stage stochastic programs for energy-resilient design, and to the risk-averse formulations of Wang and colleagues [16, 17] and Davoodi and colleagues [18]. We differ in the driving uncertainty: those studies use diffusion or scenario-tree representations, whereas we identify and use a jump process estimated from official national data. Relative to the empirical Nigerian literature [20, 22, 23, 26], which establishes the severity of the energy shock through survey and econometric evidence, our contribution is prescriptive rather than descriptive: we convert the documented shock into an investment rule with a computed shadow price. Relative to the SciML literature [5, 6, 28, 29], which reports predominantly successful applications, we contribute a well-documented failure alongside a success and a criterion for distinguishing the two cases in advance.

5.5. Contribution to the Sustainable Development Goals

The work speaks directly to SDG 7 on affordable and clean energy: the modelled transition raises the renewable share of enterprise energy consumption from zero to 72.2 per cent and eliminates unserved demand, addressing affordability and reliability together. It also speaks to SDG 3 on good health and well-being, since displacing petrol generation removes a source of localised combustion emissions and noise from workplaces that are frequently adjacent to or within residential areas.

6. Conclusion

This paper set out to determine how a Nigerian micro, small or medium enterprise should respond to energy cost shocks, using real data and defensible statistics rather than illustrative assumptions. Using the complete published National Bureau of Statistics record of national average pump prices from January 2024 to May 2026, we compared four candidate stochastic processes by maximum likelihood. A Merton jump-diffusion is selected on the Akaike criterion, its small-sample correction and the Bayesian criterion; adding jumps to a diffusion is strongly supported, including under a boundary-aware parametric bootstrap that returns a p-value of 0.01, while adding mean reversion is not. In the selected model the diffusive drift is statistically indistinguishable from zero while the jump component accounts for essentially the whole of the estimated log-price drift of 0.373 per year, and the implied probability of at least one jump in a month is 0.43. The statistically identified jump months correspond closely to the documented record of refinery entry, downstream price competition and import-duty changes.

Embedding this process in a two-stage mean-conditional-value-at-risk program with an energy-balance merit-order recourse, we find that optimal solar and storage investment reduces the representative firm's expected annual energy cost by 42.8 per cent and the conditional value-at-risk of its operating cost by 86.1 per cent, with generator dependence falling from 48.1 to 2.3 per cent of demand. Sweeping the investment budget reveals a shadow price of capital reaching about 2.67 naira of annual objective improvement per naira invested, with the cost-minimising scale delivering a 52.6 per cent reduction in expected total cost at a payback of about one year. Within the model, capital access rather than technology cost or tariff policy is the binding constraint; whether it is the binding constraint for real firms is a question their financing data would have to settle. Storage does not enter the risk-neutral optimum and appears only as risk aversion rises, identifying it as a risk-management asset.

The scientific machine learning layer produced two opposite outcomes, both reported. Learning the price drift failed: with twenty-eight transitions the unconstrained network achieved the best training error and the worst generalisation error, and the parametric process outperformed both learned alternatives out of sample. Learning the recourse operator succeeded: the surrogate reproduced expected cost and conditional value-at-risk to within 0.92 and 2.47 per cent and ran about 100 times faster. The distinguishing factor was the availability of training data, not the difficulty of the task.

Future work should extend the analysis to a panel of enterprise archetypes, replace the monthly recourse with an hourly formulation carrying explicit state-of-charge dynamics, introduce dependence between fuel prices and grid availability, incorporate equipment degradation and failure, and validate the assumed demand profile against metered enterprise data. The framework itself is portable: any economy with published monthly energy-price statistics and a large self-generating enterprise sector can apply it without firm-level data collection.

Acknowledgment

The authors received no specific grant from any funding agency in the public, commercial or not-for-profit sectors. The authors thank the National Bureau of Statistics of Nigeria and the Nigerian Electricity Regulatory Commission for maintaining the public data series on which this study depends.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this manuscript.

Data availability

All data used in this study are published official statistics; no proprietary or simulated data are used. All data underlying the results are contained within this article: the complete twenty-nine-month PMS series, the seventeen-month AGO series and the Band A tariff steps are reproduced in full in Appendix A (Table A1). The primary sources are the National Bureau of Statistics *Premium Motor Spirit (Petrol) Price Watch* and *Automotive Gas Oil (Diesel) Price Watch* [39, 40] at https://nigerianstat.gov.ng and the Nigerian Electricity Regulatory Commission Band A tariff orders. The estimation, optimization and machine-learning code that reproduces every table and figure is available from the corresponding author on reasonable request.

Funding

The authors received no external funding for this study.

References

[1] Small and Medium Enterprises Development Agency of Nigeria & National Bureau of Statistics, National survey of micro, small and medium enterprises, SMEDAN/NBS, Abuja, Nigeria, 2021. Available online: https://smedan.gov.ng/reports.

[2] African Development Bank Group, African economic outlook 2026, African Development Bank Group, Abidjan, Cote d'Ivoire, 2026. Available online: https://www.afdb.org/en/documents/african-economic-outlook-2026.

[3] World Bank, Enterprise surveys: firm-level data on infrastructure and electricity constraints, World Bank Group, Washington DC, USA, 2024. Available online: https://www.enterprisesurveys.org.

[4] Nigerian Electricity Regulatory Commission, Service-based tariff orders: Band A end-user tariff review, NERC, Abuja, Nigeria, 2024. Available online: https://nerc.gov.ng/resource-category/nerc-reports/.

[5] M. Liu, Z. O'Neill, J. Wen, T. Wu, B. Dong & Z. Yang, "Scientific machine learning (SciML) for building technology: a new paradigm for design and operations", Building and Environment **296** (2026) 114518. https://doi.org/10.1016/j.buildenv.2026.114518.

[6] H. Mohammadghasemi, J. F. Soesanto, A. Singh, B. Maciszewski & B. Huang, "Scientific machine learning for modeling industrial-scale primary separation vessel", Computers and Chemical Engineering **205** (2026) 109427. https://doi.org/10.1016/j.compchemeng.2025.109427.

[7] A. Ganczarek-Gamrot, A. Gorczyca-Goraj, K. Pilot & K. Kania, "Electricity price volatility and the performance of machine learning forecasting models in European energy markets", Energies **18** (2025) 6535. https://doi.org/10.3390/en18246535.

[8] K. Ignatieva & P. Wong, "Empirical analysis of crude oil dynamics using affine vs. non-affine jump-diffusion models", Journal of Empirical Finance **78** (2024) 101519. https://doi.org/10.1016/j.jempfin.2024.101519.

[9] W. Ben Romdhane & H. Boubaker, "A hybrid HAR-LSTM-GARCH model for forecasting volatility in energy markets", Journal of Risk and Financial Management **19** (2026) 77. https://doi.org/10.3390/jrfm19010077.

[10] B. Moghadaspoor, R. Tavakkoli-Moghaddam, A. Bozorgi-Amiri & T. Allahviranloo, "Energy-resilient closed-loop supply chain design managed by the 3PL provider: a pick-up strategy and data envelopment analysis", Journal of Industrial Information Integration **44** (2025) 100763. https://doi.org/10.1016/j.jii.2024.100763.

[11] M. Mozaffari, M. Abedini, H. Monsef & M. Marzband, "Stochastic modeling and techno-economic optimization for distribution network resilience under extreme events with integration of electric vehicles", Sustainable Energy, Grids and Networks **47** (2026) 102351. https://doi.org/10.1016/j.segan.2026.102351.

[12] M. Oszczypala, J. Konwerski, J. Ziolkowski & J. Malachowski, "Reliability analysis and redundancy optimization of k-out-of-n systems with random variable k using continuous time Markov chain and Monte Carlo simulation", Reliability Engineering and System Safety **242** (2024) 109780. https://doi.org/10.1016/j.ress.2023.109780.

[13] S. Jadhav & A. Kumar, "Stochastic modeling and availability optimization of wireless sensor network through particle swarm optimization", Reliability Engineering and System Safety **265** (2026) 111538. https://doi.org/10.1016/j.ress.2025.111538.

[14] M. B. Şenol & E. A. Murat, "A sequential solution heuristic for continuous facility layout problems", Annals of Operations Research **320** (2023) 355. https://doi.org/10.1007/s10479-022-04907-w.

[15] G. A. Adepoju, M. O. Okelola & M. A. Tijani, "Hybrid optimization technique for solving economic dispatch problem: a case study of Nigerian thermal power system", African Scientific Reports **1** (2022) 81. https://doi.org/10.46481/asr.2022.1.2.28.

[16] Y. Wang, Y. Liu & X. Bai, "Designing a new robust resilience supply chain network under partial distribution information", Computers and Industrial Engineering **190** (2024) 110028. https://doi.org/10.1016/j.cie.2024.110028.

[17] C. J. Wang & L. M. T. Zhong, "Reliable design of humanitarian supply chain under correlated disruptions: a two-stage distributionally robust approach", Annals of Operations Research **355** (2024) 2999. https://doi.org/10.1007/s10479-024-05916-7.

[18] S. M. R. Davoodi, E. Gholamian & T. Hanne, "A two-stage stochastic model considering conditional risk in a multi-level multi-product supply chain", International Journal of Optimization and Control: Theories and Applications **16** (2025) 54. https://doi.org/10.36922/IJOCTA025300131

[19] N. Chakrabarty, K. M. Sullivan & D. B. L. da Silva, "Time-based redeployment of multi-class nodes for reliable wireless sensor network coverage", Computers and Industrial Engineering **196** (2024) 110549. https://doi.org/10.1016/j.cie.2024.110549.

[20] B. I. Gumel, "The impact of fuel subsidy removal and foreign exchange harmonization on SMEs in Nigeria", SEISENSE Business Review 6 (2026) 36. https://doi.org/10.33215/tyw9t764.

[21] A. P. Priya & M. Mathew, "Business model innovation as a driver of sustainability and resilience in MSMEs: a systematic review and research agenda", Environmental and Sustainability Indicators 31 (2026) 101344. https://doi.org/10.1016/j.indic.2026.101344.

[22] D. Y. Achmad & A. P. Subriadi, "Challenges and opportunities of MSMEs in adopting blockchain technology in Industry 4.0 era", Procedia Computer Science **284** (2026) 436. https://doi.org/10.1016/j.procs.2026.07.022.

[23] P. E. Edoja, R. N. Okoh, E. O. Udueni & G. C. Aye, "Macroeconomic drivers of poultry price volatility in Nigeria: a study of inflation and exchange rate dynamics", Commodities **5** (2026) 3. https://doi.org/10.3390/commodities5010003.

[24] F. Edobor & A. Sambo-Magaji, *Small and medium enterprises (SMEs) and sustainable economic development*, in *Digital Transformation for Business Sustainability and Growth in Emerging Markets*, Emerald Publishing, Bingley, United Kingdom, 2025, pp. 197–222. https://doi.org/10.1108/978-1-83549-109-620251008.

[25] J. A. Ahile, O. C. Meludu, A. S. Oniku, S. A. Sunu, L. P. Kenda, S. Kwarki & J. O. Osumeje, "Examination of the potential for geothermal energy in parts of the Benue trough, Nigeria, through the use of high-resolution aeromagnetic data", Recent Advances in Natural Sciences **2** (2024) 124. https://doi.org/10.61298/rans.2024.2.2.124.

[26] Z. Y. I. Abba, N. Balta-Ozkan & G. Drew, "Catalysing decentralised renewable energy investment in Nigeria: investor-focused risk evaluation and de-risking strategies", Renewable and Sustainable Energy Transition **7** (2025) 100112. https://doi.org/10.1016/j.rset.2025.100112.

[27] L. C. Sim & S. Griffiths, "Renewable energy supply chains between China and the Gulf states: resilient or vulnerable?", Energy Strategy Reviews **56** (2024) 101605. https://doi.org/10.1016/j.esr.2024.101605.

[28] Y. Gong & S. H. Cheung, "Adaptive physics-informed neural operator and subset simulation for high-dimensional reliability analysis with multiple stochastic processes in stochastic differential equations", Reliability Engineering and System Safety **268** (2026) 111949. https://doi.org/10.1016/j.ress.2025.111949.

[29] J. Zhang, X. Cai, Z. Cui & J. Chen, "A physics-informed neural network surrogate model and many-objective optimization algorithm for coupled multi-energy systems in smart grids", Expert Systems with Applications **298** (2026) 129760. https://doi.org/10.1016/j.eswa.2025.129760.

[30] S. Wei, J. Jiang, J. Chen & Q. Wang, "Digital technology and supply chain resilience: a literature review and emerging themes", International Transactions in Operational Research **33** (2026) 2167. https://doi.org/10.1111/itor.70119.

[31] X. Li, V. Krivtsov, C. Pan, A. Nassehi, R. X. Gao & D. Ivanov, "End-to-end supply chain resilience management using deep learning, survival analysis, and explainable artificial intelligence", International Journal of Production Research **63** (2024) 1174. https://doi.org/10.1080/00207543.2024.2367685.

[32] A. Singh & S. B. Singh, "Dynamic reliability and sensitivity analysis of weighted k-out-of-n cold standby system with multi-performance multi-state components", Reliability Engineering and System Safety **262** (2025) 111221. https://doi.org/10.1016/j.ress.2025.111221.

[33] Y. Shen, M. Fouladirad & A. Grall, "Mathematical modeling of solar farm performance degradation in a dynamic environment for condition-based maintenance", Reliability Engineering and System Safety **257** (2025) 110778. https://doi.org/10.1016/j.ress.2024.110778.

[34] A. S. Bhandari, A. Kumar & M. Ram, "Hybrid PSO-GWO algorithm for reliability redundancy allocation problem with cold standby strategy", Quality and Reliability Engineering International **40** (2024) 115. https://doi.org/10.1002/qre.3243

[35] A. Kumar, P. Kumar & M. Forghani-elahabad, "Multistate system performance analysis incorporating human error, mathematical modeling and reliability approach", Advances in Systems Science and Applications **24** (2024) 114. https://www.researchgate.net/publication/379893356_Multistate_System_Performance_Analysis_Incorporating_Human_Error_Mathematical_modeling_and_Reliability_Approach.

[36] P. Kumar & A. Kumar, "Estimating the reliability and sensitivity of a public address system through the Markov decision process", Telecom **5** (2024) 1179.

https://doi.org/10.3390/telecom5040059.
[37] N. Patel, R. N. Rai, H. Patil & P. Shrivastava, "Reliability analysis of cutting tools for industrial applications: an integrated AHP-RSM-PHM approach", International Journal of Mathematical, Engineering and Management Sciences **9** (2024) 756. https://doi.org/10.33889/IJMEMS.2024.9.4.039.
[38] A. Kumar, V. S. Maan, R. Choudhary & M. Saini, "Availability evaluation of solar photovoltaic systems using Markov modeling and cuckoo search algorithm", Journal of Intelligent and Fuzzy Systems **46** (2024) 2261. https://doi.org/10.3233/JIFS-231940.
[39] National Bureau of Statistics, Premium motor spirit (petrol) price watch, monthly reports January 2024 to May 2026, NBS, Abuja, Nigeria, 2026. Available online: https://microdata.nigerianstat.gov.ng/index.php/catalog/157/related-materials.
[40] National Bureau of Statistics, Automotive gas oil (diesel) price watch, monthly reports 2024 to 2026, NBS, Abuja, Nigeria, 2026. Available online: https://microdata.nigerianstat.gov.ng/index.php/catalog/158/related-materials.

Appendix A. Data used in the study

Table A1: Complete monthly PMS and AGO series used in the study (NGN per litre).

Month	PMS	AGO	Month	PMS	AGO
2024-01	668.30	-	2025-04	1239.33	1722.45
2024-02	679.36	1257.06	2025-05	1027.76	1758.26
2024-03	696.79	1341.16	2025-06	1037.66	1813.81
2024-04	701.24	1415.06	2025-07	1024.99	-
2024-05	769.62	1403.96	2025-08	988.25	-
2024-06	750.17	1462.98	2025-09	970.59	1277.81
2024-07	770.54	1379.48	2025-10	1052.31	1398.57
2024-08	830.46	1406.13	2025-11	1061.35	1409.61
2024-09	1030.46	-	2025-12	1048.63	-
2024-10	1184.83	-	2026-01	1034.76	1361.57
2024-11	1214.17	1446.83	2026-02	1051.47	1420.17
2024-12	1189.12	-	2026-03	1288.54	-
2025-01	1258.34	-	2026-04	1532.93	-
2025-02	1245.80	1501.05	2026-05	1596.25	-
2025-03	1261.65	-			

PMS covers all twenty-nine months; AGO covers the seventeen verified months, with dashes where NBS did not publish a figure.

Table A1 lists the complete monthly series used for estimation and simulation. The PMS series is used for model estimation, scenario generation and all optimization; the AGO series is used only for the descriptive comparison of Section 4.1. All values are national average retail prices in naira per litre from the National Bureau of Statistics [39, 40]. The Band A end-user tariff steps used are 68.00 (to March 2024), 225.00 (April 2024), 206.80 (May 2024) and 209.50 naira per kilowatt-hour (from July 2024) [4].

Table A2: Optimal solution underlying the sensitivity analysis of Table 11. For each parameter and each low/base/high level, the table reports the optimal first-stage capacities and the optimised mean-CVaR objective (million naira), so that the re-optimisation and the monotone saving of Table 11 can be verified.

Parameter	Level	Value	y_{pv} (kWp)	y_b (kWh)	Objective (m NGN)
Cost of capital	Low	0.18	4.00	5.0	2.659
	Base	0.25	4.00	5.0	2.990
	High	0.32	3.50	3.0	3.290
Genset efficiency (L/kWh)	Low	0.34	3.50	3.0	2.867
	Base	0.42	4.00	5.0	2.990
	High	0.5	4.00	5.0	3.089
Battery cost (NGN/kWh)	Low	350,000	4.00	6.0	2.757
	Base	450,000	4.00	5.0	2.990
	High	550,000	3.50	3.0	3.115
PV capital cost (NGN/kWp)	Low	700,000	4.00	6.0	2.709
	Base	900,000	4.00	5.0	2.990
	High	1,100,000	3.50	3.0	3.229
Peak sun hours (h/day)	Low	4.4	4.00	4.0	3.103
	Base	4.9	4.00	5.0	2.990
	High	5.4	3.50	5.0	2.855
Daytime coincidence share	Low	0.45	3.50	6.0	3.364
	Base	0.55	4.00	5.0	2.990
	High	0.65	4.00	3.0	2.574
Value of lost load (NGN/kWh)	Low	1,000	4.00	5.0	2.990
	Base	1,500	4.00	5.0	2.990
	High	2,200	4.00	5.0	2.990
Band A tariff (NGN/kWh)	Low	180	4.00	5.0	2.896
	Base	209.5	4.00	5.0	2.990
	High	240	4.00	5.0	3.088
Mean grid availability	Low	0.4	4.00	5.0	2.990
	Base	0.55	4.00	5.0	2.990
	High	0.7	4.00	5.0	2.990