

LNCs 16982

Bo Nørregaard Jørgensen ·
Zheng Grace Ma · Hongbo Duan · Lei Zhu ·
Zhaohua Wang · Lin Zhang · Huibin Du ·
Tiejun Ma · Xu Tang · Qunwei Wang ·
Sanmang Wu · Yuejun Zhang (Eds.)

Energy Informatics

6th Energy Informatics Academy Conference, EI.A 2026
Beijing, China, October 14–17, 2026
Proceedings, Part II

2
Part II

EI.A 2026

Energy Informatics Academy
The Global Community for Energy Informatics

 Springer

MOREMEDIA 

Lecture Notes in Computer Science

16982

Founding Editors

Gerhard Goos
Juris Hartmanis

Editorial Board Members

Elisa Bertino , *Purdue University, West Lafayette, IN, USA*
Wen Gao, *Peking University, Beijing, China*
Bernhard Steffen , *TU Dortmund University, Dortmund, Germany*
Moti Yung , *Columbia University, New York, NY, USA*



Dynamic Tariff Design for Demand-Side Flexibility: A Regime-Switching Approach to Consumption Forecasting in Distribution Grids

Lucas Brylle^(✉)  and Henrik Madsen 

Department of Applied Mathematics and Computer Science, Technical University of Denmark, 2800 Kgs. Lyngby, Denmark
lucas@bryllesystems.com

Abstract. Dynamic tariffs require forecasts of how consumption responds to price signals not yet observed. We develop a Markov-Switching Seasonal Autoregressive model with Exogenous variables (MS-SARX) that endogenously identifies latent consumption regimes and estimates regime-dependent price elasticities via an L_2 -regularised Expectation-Maximisation algorithm. Evaluated on operational smart-meter data from Denmark's largest distribution system operator across three case studies, MS-SARX reduces root mean square error by 6–22 % at the primary substation level relative to the strongest baseline per case, with the largest gains where consumption exhibits multimodal regime structure. Embedding these forecasts in a revenue-neutral tariff optimisation yields model-implied peak-load reductions of 9–49 % and potential infrastructure deferral savings of 0.7–47.7 million DKK. These tariff and economic outcomes are simulation-based and contingent on the stability of regime-specific price response under tariffs not yet observed; field validation remains necessary. While demonstrated on a Danish DSO, the framework transfers to other tariff regimes, countries, and customer mixes wherever operational smart-meter and price data are available.

Keywords: Markov-switching models · demand-side flexibility · dynamic tariffs · distribution grid forecasting · regime-switching

1 Introduction

Accelerating electrification of transport and heating in Denmark is expected to triple distribution-grid investment needs from approximately DKK 2 billion to DKK 6.7 billion annually over the coming decade [11]. Demand-side flexibility (DSF) through dynamic tariffs offers a scalable alternative to traditional reinforcement: European studies estimate that flexibility activation could reduce distribution investment by 27–80 % [7]. Unlike explicit flexibility schemes that require bilateral contracts, implicit price-signal mechanisms scale across heterogeneous customer bases with low per-participant administrative cost [1, 7]. A

companion paper [3] develops the policy and scalability argument for implicit DSF; the present work focuses on the forecasting methodology that enables dynamic tariff design. However, designing effective dynamic tariffs requires forecasting how consumption will respond to tariff signals not yet observed, a counterfactual prediction problem. Conventional approaches assume a stationary price-demand relationship [15, 22], which may break down when industrial facilities switch between active and idle states at price thresholds, residential EV charging concentrates in low-price hours, and heat-pump operation creates weather-price interaction effects. Machine-learning models such as XGBoost capture non-linearity but offer little insight into how and when consumption responds to prices [5, 12].

We address this gap with a Markov-Switching Seasonal Autoregressive model with Exogenous variables (MS-SARX), validated on operational smart-meter data from Cerius/Radius, Denmark’s largest DSO, across three case studies. Our contribution is threefold. First, we extend regime-switching methodology from electricity price modelling [18, 29] to consumption forecasting for DSO tariff design. Second, we quantify when regime-switching adds value over baselines (XGBoost, naive persistence) across secondary and primary substation levels. Third, we close the loop from forecasting to grid management by embedding MS-SARX in a revenue-neutral tariff optimisation evaluated via NPV analysis. Figure 1 provides an overview of the applied methodology.




Fig. 1. Methodology overview. Left: forecasting pipeline (data, elasticity screening, MS-SARX estimation, 24 h evaluation). Right: tariff design loop (counterfactual simulation, NPV analysis).

2 Background

Short-term load forecasting spans classical ARIMA/SARIMAX models [15, 16, 22], gradient-boosted trees [5, 12], and deep-learning architectures [14]. While these achieve strong point-forecast accuracy, they are ill-suited to the counterfactual what-if analysis tariff design requires: deep-learning baselines in particular are excluded here because their black-box nature prevents extraction of the regime-specific price coefficients needed for counterfactual optimisation, and because $\approx 13,000$ hourly observations per substation is insufficient

for high-capacity models to outperform well-specified parametric ones. Markov-switching models, introduced by Hamilton [13] for business-cycle analysis, are widely applied to electricity price forecasting [18, 23, 29] and renewable generation [24, 25], but remain rare for consumption forecasting at the distribution-grid level, despite the natural interpretation of consumption regimes (e.g. facility on/off, EV charging active/inactive). Empirical short-run price elasticities range from -0.1 to -0.5 for residential customers [8, 17, 20], but static estimates cannot capture the state-dependent nature of demand response. While the Flexibility Function framework [19] models flexibility as a function of tariff trajectory and system state, our regime-switching approach can be viewed as a data-driven approximation where each regime implicitly encodes a distinct flexibility state.

Existing regime-switching applications in electricity thus focus on wholesale price dynamics rather than DSO-level consumption forecasting for tariff design, while the DSF literature typically assumes known, constant elasticities [7, 8]. Recent work on distribution tariff design under uncertainty [2, 28] highlights the need for demand models that capture heterogeneous price responses, but does not employ regime-switching structures. Data-driven approaches such as inverse optimisation for consumer clusters [27] estimate price response from operational data; we extend this by letting the price response itself switch between latent regimes, yielding regime-specific coefficients for counterfactual tariff evaluation.

3 Model

3.1 MS-SARX Specification

Let y_t denote consumption at hour t and u_t the exogenous electricity price signal. A latent regime variable $s_t \in \{1, \dots, K\}$ is assumed to follow a first-order Markov chain with transition matrix $\mathbf{P} = [p_{ij}]$, where $p_{ij} = P(s_t = j \mid s_{t-1} = i)$.

Prior to estimation, we apply the Yeo-Johnson power transformation [30] to stabilise variance, followed by differencing to enforce stationarity. Differencing orders d (regular) and D (seasonal, period s) are determined by ADF and KPSS tests at the 5% significance level. Denoting the Yeo-Johnson transform with parameter λ (estimated by maximum likelihood) as $\tilde{y}_t = \psi(y_t, \lambda)$, the differenced series is

$$y_t^* = (1 - L)^d (1 - L^s)^D \tilde{y}_t.$$

Within regime k , the differenced series follows an additive seasonal ARX process:

$$y_t^* = c_k + \sum_{i=1}^p \phi_{k,i} y_{t-i}^* + \sum_{m=1}^P \Phi_{k,m} y_{t-ms}^* + \sum_{j=0}^r \theta_{k,j} u_{t-j} + \varepsilon_{k,t}, \quad \varepsilon_{k,t} \sim \mathcal{N}(0, \sigma_k^2) \quad (1)$$

where $\phi_{k,i}$ and $\Phi_{k,m}$ are regime-specific non-seasonal and seasonal AR coefficients at lags i and ms respectively, c_k is the intercept, and $\theta_{k,j}$ are exogenous price coefficients up to lag r . Unlike the multiplicative SARIMA formulation, the non-seasonal and seasonal AR components enter additively without cross-product

terms, reducing the parameter count from $(p + P + pP)$ to $(p + P)$. The full model is denoted $\text{MS}(K)\text{-SARX}(p, d, r)(P, D)[s]$.

3.2 Estimation via EM

Parameters are estimated by the Expectation-Maximisation (EM) algorithm [4] with the forward-backward procedure [26] and L_2 regularisation. Initial regime assignments are obtained via Gaussian Mixture Model (GMM) clustering on consumption-price features, with silhouette scores guiding the number of components.

E-Step. Forward probabilities $\alpha_{t,k} = P(s_t = k \mid y_{1:t}^*, u_{1:t})$ are computed recursively:

$$\alpha_{t,k} = \frac{f(y_t^* \mid s_t = k, \mathcal{H}_{t-1}) \sum_i \alpha_{t-1,i} p_{ik}}{\sum_j f(y_t^* \mid s_t = j, \mathcal{H}_{t-1}) \sum_i \alpha_{t-1,i} p_{ij}} \quad (2)$$

where $\mathcal{H}_{t-1} = \{y_{1:t-1}^*, u_{1:t-1}\}$ is the information set and $f(\cdot)$ is the Gaussian emission density with regime-specific mean and variance. Backward probabilities $\beta_{t,k}$ yield smoothed probabilities $\gamma_{t,k} = P(s_t = k \mid y_{1:T}^*, u_{1:T})$ and joint smoothed probabilities $\xi_{t,i,j} = P(s_t = i, s_{t+1} = j \mid y_{1:T}^*, u_{1:T})$.

M-Step. Regime-specific parameters are updated via weighted ridge regression:

$$\hat{\beta}_k = \arg \min_{\beta} \sum_t \gamma_{t,k} (y_t^* - \mathbf{x}_t^\top \beta)^2 + \mu \|\beta\|_2^2 \quad (3)$$

where $\beta_k = (c_k, \phi_{k,1}, \dots, \phi_{k,p}, \theta_{k,0}, \dots, \theta_{k,r}, \Phi_{k,1}, \dots, \Phi_{k,P})^\top$ collects all regime-specific coefficients and $\mathbf{x}_t = (1, y_{t-1}^*, \dots, y_{t-p}^*, u_t, \dots, u_{t-r}, y_{t-s}^*, \dots, y_{t-P_s}^*)^\top$ stacks the corresponding regressors. Regime variances and transition probabilities update as

$$\hat{\sigma}_k^2 = \frac{\sum_t \gamma_{t,k} (y_t^* - \hat{y}_{t|k}^*)^2}{\sum_t \gamma_{t,k}}, \quad \hat{p}_{ij} = \frac{\sum_{t=2}^T \xi_{t-1,i,j}}{\sum_{t=2}^T \gamma_{t-1,i}}.$$

To prevent regime absorption, we regularise the transition matrix via

$$\mathbf{P}_{\text{reg}} = (1 - \eta) \mathbf{P} + \eta \mathbf{U}, \quad \eta = 0.10, \quad \mathbf{U} = \frac{1}{K} \mathbf{1}\mathbf{1}^\top,$$

where $\mathbf{U}$ is the uniform matrix, ensuring every state retains a minimum transition probability to all others. Since EM guarantees only local convergence [4], we restart from three GMM seeds and select the run with the highest log-likelihood. Convergence is declared when the relative log-likelihood change falls below 10^{-4} . The L_2 penalty is fixed at $\mu = 1.0$ across all cases. Although the cases differ greatly in consumption magnitude, a common penalty is justified because the ridge term in Eq. (3) acts on the Yeo-Johnson-transformed and differenced series rather than on raw kWh, which substantially reduces the scale disparity between cases and makes μ approximately comparable across them. Model complexity is in any case governed chiefly by the BIC-selected order grid rather than by μ ; systematic per-case tuning of the penalty is identified as future work (Sect. 6).

3.3 Forecasting

For a forecast horizon H , regime probabilities are propagated via $\boldsymbol{\pi}_{t+h} = \mathbf{P}^\top \boldsymbol{\pi}_{t+h-1}$, starting from filtered probabilities at time t . Point forecasts are regime-probability-weighted:

$$\hat{y}_{t+h}^* = \sum_{k=1}^K \pi_{t+h,k} \hat{y}_{t+h|k}^* \quad (4)$$

Forecasting is recursive: for $h > 1$, the AR terms in Eq. (1) that depend on not-yet-observed consumption values are replaced by their predictions $\hat{y}_{t+h-1}^*$, so forecast errors accumulate over the 24-step horizon. Original-scale predictions are recovered by inverting the differencing and Yeo-Johnson transform. All metrics are reported on original-scale consumption.

3.4 Model Selection

Hyperparameters (p, P, r, K) are selected via grid search minimising the Bayesian Information Criterion:

$$\text{BIC} = n_p \log(n) - 2\mathcal{L} \quad (5)$$

where n is the number of observations and n_p counts all regime-specific coefficients, transition probabilities, and initial-state probabilities. BIC is preferred over AIC because parameters grow multiplicatively with K , and the stronger penalty ($\log n$ vs. 2) guards against overfitting. The grid spans $K \in \{1, 2, 3\}$, $p \in \{1, \dots, 16\}$, $P \in \{1, \dots, 16\}$, $r \in \{0, \dots, 24\}$, $s \in \{5, 6, 12, 24\}$. Performance is assessed on expanding-window cross-validation: the initial training window covers approximately 80% of the data (≈ 432 days) and is expanded by one 24-hour block at each step. Each step produces a 24-hour-ahead recursive forecast evaluated on original-scale consumption (RMSE, MAE). BIC model selection is performed on the training partition only, and the evaluation windows are reserved for out-of-sample assessment. Baselines include naive persistence (24 h lag) and XGBoost configured as a multi-output regressor (learning rate 0.1, 100 boosting rounds, max depth 6, subsample ratio 0.8, column subsample ratio 0.8) with temporal features (hour-of-day, day-of-week) and lagged consumption/price at lags 1–24.

3.5 Tariff Optimisation

The forecasting model feeds into a tariff optimisation that minimises overloading at constrained substations. The tariff $\tau = (\tau_1, \dots, \tau_{24})$ is parameterised in B time bins (typically $B=3$, covering roughly 8-hour blocks) and solved via L-BFGS-B with box constraints $\tau_{\min} \leq \tau_b \leq \tau_{\max}$. The objective is a weighted penalty function:

$$\tau^* = \arg \min_{\tau} w_{\text{cap}} \sum_t [\max(0, \hat{y}_t(\tau) - C_{\text{th}})]^2 + w_{\text{rev}} \Delta R(\tau)^2 \quad (6)$$

where C_{th} is the substation capacity threshold (e.g. 7,000 kW for Case 1, 200 kVA $\approx$ 200 kW at unity power factor for Case 2), w_{cap} and w_{rev} are penalty weights, and

$$\Delta R(\tau) = \frac{\sum_t \tau_t \hat{y}_t(\tau) - R_0}{R_0}$$

is the relative revenue deviation from the baseline $R_0 = \sum_t \tau_t^{\text{base}} y_t^{\text{hist}}$. Here $\hat{y}_t(\tau)$ denotes the MS-SARX counterfactual forecast under tariff scenario τ : the total price signal $u_t = p_t^{\text{spot}} + \tau_t$ is substituted into the regime-specific price coefficients $\theta_{k,j}$ from Eq. (1), and regime probabilities are propagated forward under the new price to predict consumption. Weights $w_{\text{cap}} = 1,000$ and $w_{\text{rev}} = 100$ are tuned to reduce capacity exceedance while keeping revenue deviation within $\pm 5\%$. The revenue penalty enforces approximate revenue neutrality relative to the current Tariff Model 3.0 [9].

4 Data and Experimental Setup

4.1 Data Sources

We use operational data from Cerius/Radius, Denmark’s largest DSO, serving 1.5 million customers across Zealand. The dataset comprises: (i) hourly point-of-delivery (POD) consumption from smart meters (Oct 2023–Apr 2025, $\approx$ 13,000 observations per POD), (ii) day-ahead spot prices from Nord Pool (DK2 zone) via Energinet’s data service, and (iii) grid topology and transformer capacity data from SAP/GIS systems.

The current specification uses only price as an exogenous covariate to isolate price–consumption dynamics. Omitting temperature is a limitation for heat-pump-dominated loads (Sect. 6).

4.2 Price Elasticity Screening

To identify flexibility candidates, we estimate price elasticity for industrial customers averaging > 25 kWh/h via log-log OLS regression:

$$\ln Q_t = \alpha + \hat{e}_p \ln P_t + \mathbf{X}_t' \boldsymbol{\delta} + \nu_t,$$

where $\mathbf{X}_t$ contains temporal controls (hour-of-day, day-of-week) and ν_t is the regression residual. Customers are classified as highly elastic ($\hat{e}_p \leq -0.25$), moderately elastic ($-0.25 < \hat{e}_p \leq -0.1$), or inelastic ($\hat{e}_p > -0.1$), filtering at $p < 0.05$.

Figure 2 shows that while 94.5% of industrial customers are inelastic by count, elastic customers account for 12.9% of total industrial consumption volume, indicating concentrated flexibility potential among larger loads.

4.3 Case Study Design

Three case studies capture diverse grid challenges (Table 1):




Fig. 2. Price elasticity distribution among industrial customers. Left: by customer count. Right: by consumption volume (MWh). Elastic customers ($\hat{e}_p \leq -0.1$) represent 5.5 % of customers but 12.9 % of consumption.

Case 1: Industrial Facility. A primary substation approaching reserve capacity limits, with a major industrial customer exhibiting binary consumption (0 vs. 200–300 kW) driven by spot-price thresholds. At the individual facility level the price–consumption correlation reaches $\rho = -0.92$, dropping to $\rho = -0.78$ at the secondary substation aggregate due to co-located loads.

Case 2: Residential EV Area. A secondary substation serving 100–200 residential connections where nighttime EV charging exceeds the 200 kVA transformer rating by 20–50 % during low-price hours. Because this area contains distributed rooftop photovoltaics, the metered quantity is *net* consumption (gross load minus local generation and export) and is therefore negative during sunny, low-demand hours. This does not impair regime identification: the EM step operates on the differenced, Yeo–Johnson-transformed series, so negative net values simply populate the low-consumption regime. Price coefficients are accordingly interpreted as the response of net load—a negative $\theta_{k,j}$ reflects the combined effect of demand reduction and export increase under high prices.

Case 3: Rural Distribution. A primary substation in a rural area with long radial feeders serving a mixed customer base of small industry, residential, and seasonal vacation properties, where peak demand periodically exceeds transformer capacity.

Table 1. MS-SARX model configurations by case study. Ratio denotes training samples per parameter.

Case	Level	Specification	Params	Samples	Ratio
1	Secondary	MS(2)-SARX(11,1,12)(12,1)[5]	78	11 021	141:1
1	Primary	MS(3)-SARX(8,1,12)(6,1)[24]	93	12 672	136:1
2	Secondary	MS(1)-SARX(7,0,8)(12,1)[24]	30	11 016	367:1
2	Primary	MS(2)-SARX(16,1,24)(12,1)[6]	136	10 226	75:1
3	Secondary	MS(2)-SARX(9,0,4)(8,1)[12]	46	9 768	212:1
3	Primary	MS(2)-SARX(12,1,24)(16,1)[12]	112	10 416	93:1

BIC selection yields case-specific seasonal periods ($s=5$ for Case 1 secondary, a five-hour production cycle; $s=6$ for Case 2 primary, a six-hour EV rhythm), and selects $K=1$ for Case 2 secondary—reducing to a standard SARX—while two regimes emerge at its primary level (Sect. 5.1). The seasonal differencing order is constrained to $D \leq 1$, since higher orders cause error accumulation during the multi-step inverse transform that destabilises 24-hour-ahead forecasts.

5 Results

5.1 Regime Identification

The MS-SARX models identify interpretable consumption regimes whose transitions align with price dynamics. Figure 3 illustrates this for Case 1 at two aggregation levels over the same representative week.

At the secondary (individual facility) level, two regimes capture the binary on/off pattern of the dominant industrial customer: Regime 1 (83.2% coverage) spans the broad operating range while Regime 2 (16.8%) isolates the low-price, high-consumption state. The transition matrix exhibits high persistence: self-transition probabilities of 96.4% (Regime 1) and 64.4% (Regime 2). At the primary (aggregated) level, three regimes emerge: a baseline high-consumption, high-price state (65.0%), a low-consumption, low-price state (27.6%), and a high-consumption, low-price state (7.3%). The third regime, absent at the facility level, shows how regime structure becomes richer through aggregation; kernel density estimates confirm that Regime 2 concentrates at high consumption and low price, consistent with a facility activating below its marginal cost threshold. More broadly, the recovered regimes admit direct physical readings: in Case 1 they track the dominant facility’s on/off production states; in Case 2 the secondary regime that emerges at the primary level aligns with hours of concentrated nighttime EV charging; and in heat-pump-dominated feeders the same mechanism would be expected to separate thermostatically-driven heating bursts from baseline demand. This correspondence between latent regimes and identifiable operating states is what lets the regime-specific price coefficients be read as distinct flexibility states rather than statistical artefacts.

For Case 2, BIC selects $K=1$ at the secondary level, correctly reducing to standard SARX for this unimodal load, while at the primary level two regimes emerge as aggregation creates systematic price-driven patterns—further evidence that aggregation amplifies regime structure.

5.2 Forecasting Performance

Secondary Substation Level. Table 2 presents forecasting performance at the secondary (local) level. MS-SARX achieves the lowest RMSE in Cases 1 and 3.

For Case 1, MS-SARX improves RMSE by 3.2% over XGBoost (the naive baseline attains lower MAE through diurnal persistence but higher RMSE variability). Case 2 is near-equivalent, consistent with $K=1$ selection, while Case 3 shows the clearest advantage—an 18.3% RMSE improvement, as its mixed customer base produces regime-dependent behaviour the switching structure exploits.




Fig. 3. MS-SARX forecast for Case 1 at secondary (top, 2 regimes) and primary level (bottom, 3 regimes). Actual consumption (black), prediction (dashed red) with regime-probability shading; spot price below. (Color figure online)

Primary Substation Level. Table 3 reports primary-level performance, where MS-SARX shows consistent and substantial improvements.

Relative to the strongest baseline per case, RMSE reductions are 6.0 % in Case 1 (vs. Naive), 17.1 % in Case 2 (vs. Naive), and 22.3 % in Case 3 (vs. XGBoost). One asymmetry warrants comment: in Case 3, MS-SARX is evaluated on 82 daily forecast origins while the baselines are tabulated over the denser hourly grid (1,968 origins) for computational economy. Two considerations make the 22.3 % comparison defensible. First, the Diebold–Mariano test underlying the significance claim pairs forecast errors at common origins and is therefore unaffected by the differing window *counts*. Second, because each origin contributes a full 24-hour recursive forecast, restricting the baselines to the same 82 daily origins changes their aggregate RMSE only marginally—for the persistence baseline by under one percent—so the ranking is preserved. We nonetheless recommend reporting all models on a common daily-origin grid in any deployment-grade evaluation. In Case 1, naive persistence (RMSE 1555) outperforms XGBoost (2022) owing to strong autoregressive structure, and MS-SARX improves further by capturing regime transitions misaligned with the 24 h

Table 2. Secondary substation forecasting performance (mean $\pm$ s.d. across expanding-window evaluations; $\approx 1,900$ windows per case). Best value per case and metric in bold.

Case	Model	MAE (kWh)	RMSE (kWh)
1	MS-SARX	76.7 $\pm$ 21.2	89.8 $\pm$ 22.9
	XGBoost	78.2 $\pm$ 23.9	92.8 $\pm$ 25.6
	Naive 24 h	65.2 $\pm$ 34.3	93.6 $\pm$ 43.1
2	MS-SARX	16.6 $\pm$ 4.7	20.7 $\pm$ 5.7
	XGBoost	15.7 $\pm$ 4.8	20.9 $\pm$ 6.1
	Naive 24 h	19.2 $\pm$ 6.6	24.6 $\pm$ 8.5
3	MS-SARX	11.0 $\pm$ 5.7	13.4 $\pm$ 6.5
	XGBoost	12.2 $\pm$ 6.6	16.4 $\pm$ 7.1
	Naive 24 h	12.6 $\pm$ 7.8	15.9 $\pm$ 8.2

Table 3. Primary substation forecasting performance ($\approx 1,968$ evaluation windows; Case 3 MS-SARX: 82 windows). Best RMSE per case in bold.

Case	Model	MAE (kWh)	RMSE (kWh)
1	MS-SARX	1212 $\pm$ 498	1461 $\pm$ 568
	XGBoost	1718 $\pm$ 530	2022 $\pm$ 530
	Naive 24 h	1260 $\pm$ 449	1555 $\pm$ 518
2	MS-SARX	1586 $\pm$ 786	1952 $\pm$ 924
	XGBoost	2052 $\pm$ 811	2444 $\pm$ 1025
	Naive 24 h	1912 $\pm$ 1162	2353 $\pm$ 1392
3	MS-SARX	1320 $\pm$ 778	1580 $\pm$ 886
	XGBoost	1761 $\pm$ 792	2033 $\pm$ 837
	Naive 24 h	1977 $\pm$ 1222	2297 $\pm$ 1297

cycle. The consistently stronger primary-level gains reflect aggregation effects (Sect. 6).

Explanatory Power. To complement the error metrics, we compute the out-of-sample coefficient of determination R^2 between predicted and observed consumption across all 24 h windows. MS-SARX attains R^2 of 0.22–0.58 at the primary level and 0.25–0.69 at the secondary level—highest in Case 2 (0.69 secondary, 0.58 primary), consistent with its diurnal regularity, and lowest in Case 1 (0.25, 0.22), reflecting the facility’s irregular load and a 64 % mean consumption increase in the test period relative to training.

Statistical Significance. Diebold–Mariano tests [6] confirm that all primary-level improvements are significant ($p < 0.05$, MS-SARX vs. XGBoost for all cases; vs. Naive for Cases 2–3). At the secondary level, Cases 1 and 2 are not significant ($p > 0.10$), consistent with the modest RMSE gaps.

5.3 Tariff Optimisation

We now evaluate whether these forecasting improvements translate into operationally useful tariff designs. The peak reductions and savings reported below are *model-implied*: they are obtained by simulating consumption under counterfactual tariffs with the estimated regime-specific price coefficients, and should be read as planning estimates contingent on the stability of that price response, not as validated field effects. All tariffs use the bounds $[\tau_{\min}, \tau_{\max}] = [-0.2, 1.2]$ DKK/kWh, chosen to span the range of surcharges and rebates plausible under Tariff Model 3.0 [9].

Case 1: Industrial Capacity Constraint. Fig. 4 illustrates the optimised tariff at the primary level. The tariff raises prices to 1.2 DKK/kWh during constrained morning hours (07:00–15:00) and drops to -0.2 DKK/kWh off-peak. The industrial facility contributes a 90.3 kW peak reduction (48.9% of reserve capacity exceedance). Load is shifted rather than curtailed, consistent with revenue neutrality.




Fig. 4. Tariff optimisation for Case 1 (industrial, primary level). (a) The optimised tariff (solid red) raises peak-hour prices (07:00–15:00) to shift consumption below the 7, (Color figure online)000 kW threshold (dotted orange). (b) Consumption deviation (optimised – initial): up to 90 kW reduction during constrained hours, offset off-peak consistent with revenue neutrality.

Case 2: Residential EV Charging. Fig. 5 shows the Case 2 tariff addressing night-time EV charging peaks. High tariffs during 00:00–06:00 reduce charging peaks by 20–31 kWh (9–13%), with compensating daytime increases. As EV penetration grows, these reductions are expected to increase substantially.




Fig. 5. Tariff optimisation for Case 2 (residential EV, secondary level). (a) The optimised tariff raises nighttime prices (00:00–06:00) to shift EV charging below the 200 kWh threshold (dotted orange). (b) Consumption deviation: 20–31 kWh reductions during constrained hours, offset during daytime consistent with revenue-neutral shifting. (Color figure online)

Case 3: Rural Aggregation. Case 3 illustrates a different mechanism: no single customer provides sufficient flexibility, but the tariff acts across all secondary substations simultaneously—each contributing a modest $\sim 4.5\%$ load shift that aggregates to a cumulative 8.5 kW peak reduction per feeder. This aggregation effect, where many individually insufficient responses combine into grid-relevant flexibility, is the mechanism that makes implicit DSF scalable and where regime-switching forecasting (22.3 % RMSE improvement) provides the largest gains.

Load-Profile Quality and Bound Sensitivity. Peak reduction alone does not fully characterise the optimised profile. Because the objective is revenue-neutral, the tariff *shifts* rather than curtails load: total energy is conserved within the $\pm 5\%$ revenue band, so demand displaced from constrained hours reappears off-peak (Figs. 4b, 5b). Clipping the peak and filling the off-peak trough therefore raises the load factor and lowers the peak-to-average ratio rather than merely relocating the peak, while the three-bin parameterisation confines additional ramping to the bin boundaries; per-case values of these metrics are available from the authors. On the bounds, the optimiser saturates $[\tau_{\min}, \tau_{\max}]$ during constrained hours because the capacity constraint is binding: within revenue neutrality the objective favours the maximal admissible price spread. The reported reductions are thus an *upper envelope* for these bounds; narrower bounds produce proportionally smaller but qualitatively identical shifts, leaving the relative ranking across cases unchanged.

Economic Viability. Table 4 compares flexibility against traditional reinforcement via NPV analysis (5 % discount, 3–8 year horizons), using Cerius/Radius internal investment estimates per substation; “Flex Only” assumes indefinite deferral and “Hybrid” eventual reinforcement.

Table 4. Net present value analysis (million DKK). “Hybrid” denotes flexibility with eventual deferred reinforcement. Savings range: lower bound is Traditional – Hybrid (deferred reinforcement); upper bound is Traditional – Flex Only (indefinite deferral).

Case	Traditional (M DKK)	Flex Only (M DKK)	Hybrid (M DKK)	Savings (M DKK)
1	−10.00	−0.65	−9.29	0.71–9.35
2	−0.75	−0.05	−0.70	0.05–0.70
3	−50.00	−2.32	−36.16	13.84–47.68

6 Discussion

When Does Regime-Switching Add Value? MS-SARX improvements are largest where consumption is multimodal—industrial on/off states (Case 1), concentrated EV charging (Case 2 primary), and heterogeneous customer mixes (Case 3)—and the model correctly reduces to $K = 1$ when consumption is unimodal (Case 2 secondary). The stronger primary-level gains (6–22 %) reflect aggregation: diverse price-responsive behaviours create clearer regime structure as individual noise averages out, precisely where capacity constraints drive major DSO investment.

Operationalising the Framework. For a DSO the pipeline slots into existing planning cycles. For each constrained substation, the regime-specific elasticities and counterfactual optimiser produce a tariff schedule and an expected load-shifting envelope that can feed the periodic grid-investment review, quantifying how much reinforcement a price signal might defer before capital is committed. The same regime diagnostics support regulatory dialogue under Tariff Model 3.0 [9], where revenue neutrality and equity must be demonstrated to the regulator. Most directly, it is a screening and design tool for field pilots, identifying which substations and segments carry exploitable flexibility (Sect. 4.2) and proposing the specific tariff to trial.

Counterfactual Validity and Circularity. The price coefficients $\theta_{k,j}$ are estimated under the existing static tariff and may not transfer to a dynamic setting—the Lucas critique [21]. The regime-switching structure mitigates this: decomposing

a globally nonlinear price–consumption relationship into piecewise-linear segments limits extrapolation to within-regime neighbourhoods, whereas a single-regime model extrapolates along one global gradient mixing distinct states. MS-SARX instead activates the regime whose historical price range best matches the counterfactual scenario. It nonetheless assumes regime coefficients remain stable under moderate tariff perturbations, and using one model for both optimisation and evaluation creates a circularity that likely overstates peak-reduction effectiveness; field trials are needed to quantify this. On equity, revenue neutrality preserves aggregate charges but shifts costs in time, so inelastic customers may face higher peak-hour charges—Tariff Model 3.0 [9] provides equity review provisions for deployment.

Consumer Contract Structure. The simulation assumes consumers respond to $u_t = p_t^{\text{spot}} + \tau_t$, whereas some residential customers remain on fixed-price contracts—a share declining as Market Model 3.0 [10] promotes spot-price exposure. Our results thus apply primarily to spot-exposed customers (industrial and large commercial), which already cover most of the flexibility potential identified in Sect. 4.2, and the framework broadens as residential spot adoption grows.

Modelling Framework Under Future Tariff Regimes. The ARX specification suits the current setting: under Tariff Model 3.0 [9] tariffs are set independently of real-time consumption, so the price signal u_t enters legitimately as an exogenous covariate. Under future dynamic tariffs, however, prices become functions of aggregate consumption, breaking this exogeneity; capturing the resulting consumption–tariff feedback then requires extending MS-SARX to a multivariate regime-switching specification [22] before deployment.

Limitations and Future Work. The OLS elasticity screening does not address price–demand simultaneity [17], and the first-order Markov assumption may not capture seasonal regime prevalence. Excluding temperature means that price coefficients $\theta_{k,j}$ partially absorb weather-driven patterns that co-vary with spot prices through the merit order. Incorporating weather covariates (e.g. temperature) into $\mathbf{x}_t$ would separate genuine price response from seasonal demand—a promising direction we expect to refine rather than overturn the present conclusions, since the flexibility-relevant loads here are price- rather than weather-driven. Future extensions include probabilistic Flexibility Functions [19], graph-based voltage constraints, hybrid regime-switching/neural architectures, and systematic tuning of μ .

7 Conclusion

We presented MS-SARX, a regime-switching forecasting framework for demand-side flexibility through dynamic tariffs in distribution grids. It reduces primary-level forecasting error by 6–22% relative to the strongest baseline per case (Diebold–Mariano $p < 0.05$), with the largest gains where consumption is

multimodal—and is most valuable at the primary substation level, where capacity constraints drive investment and aggregation sharpens regime structure. Embedding these forecasts in revenue-neutral tariff optimisation yields model-implied peak-load reductions of 9–49% and infrastructure deferral savings of 0.7–47.7 million DKK; these figures are simulation-based and contingent on the stability of price response under new tariffs. The framework is not specific to Denmark—it needs only operational smart-meter consumption and a price signal, transferring to other countries, tariff regimes, and customer mixes, though regimes and elasticities are local and must be re-estimated per network. Field trials across consumption segments are needed to validate these simulated gains under real dynamic tariff deployment.

Acknowledgments. This work is supported by INSIEME (Digital Europe No. 101194952), ELEXIA (Horizon Europe No. 101075656), ARV (EU H2020 No. 101036723), SEEDS (Horizon Europe No. 10042024), and IEA EBC Annex 96 (EUDP No. 134251-549133). The authors sincerely thank Niels Andersen (Cerius/Radius) for his valuable domain guidance and support throughout this work, and gratefully acknowledge Cerius/Radius for providing access to operational grid data under data-protection agreements ensuring customer anonymity.

Disclosure of Interests. The operational data were provided by Cerius/Radius under a data-protection agreement. The authors declare no competing interests.

Data Availability. The Cerius/Radius smart-meter and grid data are confidential and not publicly available owing to customer-privacy and commercial restrictions; aggregated, anonymised results are reported in full in the paper, and the model implementation is available from the corresponding author on reasonable request.

References

1. Askeland, M., Backe, S., Bjarghov, S., Lindberg, K.B., Korpås, M.: Activating the potential of decentralized flexibility and energy resources to increase the EV hosting capacity: a case study of a multi-stakeholder local electricity system in Norway. *Smart Energy* **3**, 100034 (2021). <https://doi.org/10.1016/j.segy.2021.100034>
2. Askeland, M., Burandt, T., Gabriel, S.A.: A stochastic MPEC approach for grid tariff design with demand-side flexibility. *Energy Syst.* **14**, 707–729 (2023). <https://doi.org/10.1007/s12667-020-00407-7>
3. Brylle, L., Madsen, H., Andersen, N.: Scaling demand-side flexibility through dynamic tariffs. Submitted to EIA Nordic 2026 (2026). submitted
4. Dempster, A.P., Laird, N.M., Rubin, D.B.: Maximum likelihood from incomplete data via the EM algorithm. *J. Roy. Stat. Soc. B* **39**(1), 1–38 (1977)
5. Diaz-Iglesias, A., Belaunzaran, X., Florez-Tapia, A.M.: Short-term power demand forecasting for diverse consumer types to enhance grid planning and synchronisation. arXiv preprint [arXiv:2506.04294](https://arxiv.org/abs/2506.04294) (2025)
6. Diebold, F.X., Mariano, R.S.: Comparing predictive accuracy. *J. Bus. Econ. Stat.* **13**(3), 253–263 (1995)
7. DNV: Demand-side flexibility in the EU: quantification of benefits in 2030. Tech. rep., DNV for smartEn (Smart Energy Europe) (2022). https://smarten.eu/wp-content/uploads/2022/10/SmartEN-DSF-benefits-2030-Report_DIGITAL-1.pdf

8. Faruqui, A., Sergici, S.: Household response to dynamic pricing of electricity: a survey of 15 experiments. *J. Regul. Econ.* **38**(2), 193–225 (2010)
9. Forsyningstilsynet: Tarifmodel 3.0: regulation of network tariffs in the electricity sector. Tech. rep., Forsyningstilsynet, Copenhagen (2024)
10. Gade, P.A.V., Skjøtskift, T., Bindner, H.W., Kazempour, J.: Ecosystem for demand-side flexibility revisited: the Danish solution. *Electr. J.* **35**(9), 107206 (2022). <https://doi.org/10.1016/j.tej.2022.107206>
11. Green Power Denmark: Elnet til meget mere: investeringer i eldistributionssystemet 2024–2040 og implementering af grønt tillæg. Tech. rep., Green Power Denmark (2024). <https://greenpowerdenmark.dk/files/media/document/Elnet-til-meget-mere-analyse.pdf>, in Danish. English summary: Electricity grids for much more—investments in the distribution grid 2024–2040
12. Haben, S., Giasemidis, G., Ziel, F., Arora, S.: Short-term load forecasting and the effect of temperature at the low voltage level. *Int. J. Forecast.* **35**(4), 1469–1484 (2019)
13. Hamilton, J.D.: A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica* **57**(2), 357–384 (1989)
14. Hertel, M., Ott, S., Neumann, O., Schäfer, B., Mikut, R., Hagenmeyer, V.: Transformer neural networks for building load forecasting. In: Tackling Climate Change with Machine Learning, Workshop at NeurIPS 2022 (2022). <https://www.climatechange.ai/papers/neurips2022/43>
15. Hong, T., Fan, S.: Probabilistic electric load forecasting: a tutorial review. *Int. J. Forecast.* **32**(3), 914–938 (2016)
16. Hyndman, R.J., Athanasopoulos, G.: Forecasting: Principles and Practice, 3rd edn. OTexts (2021)
17. Ito, K.: Do consumers respond to marginal or average price? Evidence from non-linear electricity pricing. *Am. Econ. Rev.* **104**(2), 537–563 (2014)
18. Janczura, J., Weron, R.: An empirical comparison of alternate regime-switching models for electricity spot prices. *Energy Econ.* **32**(5), 1059–1073 (2010)
19. Junker, R.G., et al.: Characterizing the energy flexibility of buildings and districts. *Appl. Energy* **225**, 175–182 (2018)
20. Lijesen, M.G.: The real-time price elasticity of electricity. *Energy Econ.* **29**(2), 249–258 (2007)
21. Lucas, R.E., Jr.: Econometric policy evaluation: a critique. *Carn.-Roch. Conf. Ser. Public Policy* **1**, 19–46 (1976)
22. Madsen, H.: Time Series Analysis. Chapman & Hall/CRC (2007)
23. Mount, T.D., Ning, Y., Cai, X.: Predicting price spikes in electricity markets using a regime-switching model with time-varying parameters. *Energy Econ.* **28**(1), 62–80 (2006)
24. Pinson, P., Christensen, L.E.A., Madsen, H., Sørensen, P.E., Donovan, M.H., Jensen, L.E.: Regime-switching modelling of the fluctuations of offshore wind generation. *J. Wind Eng. Ind. Aerodyn.* **96**(12), 2327–2347 (2008)
25. Pinson, P., Madsen, H.: Adaptive modelling and forecasting of offshore wind power fluctuations with Markov-switching autoregressive models. *J. Forecast.* **31**(4), 281–313 (2012)
26. Rabiner, L.R.: A tutorial on hidden Markov models and selected applications in speech recognition. *Proc. IEEE* **77**(2), 257–286 (1989)
27. Saez-Gallego, J., Morales, J.M., Zugno, M., Madsen, H.: A data-driven bidding model for a cluster of price-responsive consumers of electricity. *IEEE Trans. Power Syst.* **31**(6), 5001–5011 (2016)

28. Schittekatte, T., Meeus, L.: Least-cost distribution network tariff design in theory and practice. *Energy J.* **41**(5), 119–156 (2020). <https://doi.org/10.5547/01956574.41.5.tsch>
29. Weron, R.: Electricity price forecasting: a review of the state-of-the-art with a look into the future. *Int. J. Forecast.* **30**(4), 1030–1081 (2014)
30. Yeo, I.K., Johnson, R.A.: A new family of power transformations to improve normality or symmetry. *Biometrika* **87**(4), 954–959 (2000)