Holding the Line: A Digital Twin Framework Supporting Policy Design for U.S. Power Grid Resilience

Daniel Freedman, Alison Pike, Avery Suriano, Walker Watson, and Henning Mortveit.*
**Corresponding Author: Henning.Mortveit@virginia.edu*

Department of Systems and Information Engineering
University of Virginia
Charlottesville, Virginia, USA

***Abstract*—The U.S. power grid faces growing reliability challenges as rising peak demand, renewable intermittency, and baseload capacity retirements strain existing infrastructure. While long-term solutions require grid expansion, near-term resilience depends on optimizing current resources. Price-based demand response (PBDR) incentivizes households to shift electricity use away from peak periods; however, existing evaluations aggregate load reductions across entire service territories, obscuring impacts on specific grid assets and communities. To address this gap, this paper develops:**

- **A high-resolution digital twin framework that couples a synthetic population, hourly electricity demand profiles, and a distribution network to derive system-level demand reductions from individual households**
- **A stochastic adoption model and demand reallocation algorithm to capture how household demographics drive load shifts across the distribution network**
- **A case study of Virginia's Eastern Shore under two candidate PBDR policies, producing 4–5% peak load reductions consistent with empirical literature while resolving shifts at the household and substation levels**

The results demonstrate that established demand response effects can be reproduced at household and substation resolution without loss of fidelity. By linking demographic characteristics and load-shifting behavior to substation-level outcomes, the framework can inform more resilient demand-side energy policy. This framework is extendable nationwide, providing a scalable approach to evaluate the grid impacts of utility policies.

***Index Terms*—digital twin, demand response, electrical distribution networks, Monte Carlo simulation, Time-of-Use pricing, peak load reduction, synthetic population modeling**

I. Introduction

The United States power grid faces mounting stress. After decades of average annual load growth below 1%, peak demand is rising, driven by electrification, artificial intelligence, and data centers [1]. Simultaneous retirements of legacy generation plants are eroding capacity and reserves, transmission constraints are creating congestion on the grid, and variable renewable generation introduces intermittency. Together, these factors strain grid reliability. New generation and transmission infrastructure is needed, but only 19% of projects requesting interconnection reach commercial operations, with average project completion times growing to 55 months in 2024 [2].

Price-based demand response (PBDR) offers a more immediate lever. Time-varying electricity prices encourage consumers to shift non-essential consumption away from peak periods; empirical studies of Time-of-Use (TOU) programs report peak reductions of 5-15% [3]. Yet participation and responsiveness vary between demographic groups, and most PBDR evaluations are modeled at the system or feeder level, obscuring the substation-scale outcomes that matter for distribution planning [4].

This paper addresses that gap. We construct a high-resolution digital twin of Virginia's Eastern Shore (Accomack and Northampton counties as shown in Figure 1) comprising a synthetic household population, annual hourly load profiles at the household level, and a synthetic distribution network linking residences to substations. A logistic adoption model assigns participation probabilities as an empirical function of household age, income, and occupancy. Monte Carlo simulation propagates behavioral uncertainty to substation-level load. We evaluate two PBDR regimes (TOU tariffs and French Tempo-style tiered-day tariffs), measuring peak load reductions, hourly demand shifts, and demographic variation in participation. The framework is designed to be extendable to other regions and stress scenarios. The household-level resolution of the digital twin is intentional, as it reflects the level at which PBDR policy is actually formulated and adopted. Unlike aggregate models, which must assume how demographic variation translates to load response and offer limited ways to validate results at the substation scale, this framework makes both the policy mapping and its demographic effects transparent and testable. In doing so, it provides a flexible framework that policymakers can use to evaluate demand response strategies within their own regional contexts.

This work was partially funded by NSF #2053013.

II. Background

A. Electric Grid Structure

Electric power systems comprise high-voltage transmission networks that transport electricity from generators over long distances and lower-voltage distribution networks that deliver it to end users. Substations interface the two tiers, stepping down voltage and feeding localized service areas. Infrastructure must be sized to peak load, meaning maximum demand over a given interval, rather than average demand, making peak conditions the critical constraint in system planning [5].

B. Price-Based Demand Response (PBDR)

Demand response programs alter electricity consumption in response to external signals and are broadly classified as incentive-based (compensating users for load reductions) or price-based (exposing consumers to time-varying rates). PBDR rests on the principle of price elasticity, assuming consumer behavior is sensitive to price. Common structures include Time-of-Use (TOU) pricing, which defines fixed peak and off-peak windows within a day, and Tempo pricing, which differentiates rates across days of the month, designating each day as one of three tariff levels based on anticipated system demand [6], [7].

C. Modeling Approaches

Bottom-up demand models construct aggregate load from individual end uses such as heating, cooling, and appliances. Digital twin frameworks extend this approach by embedding synthetic consumers within a spatially explicit distribution network, enabling substation-level analysis of household behavior. Because residential electricity demand is inherently stochastic, Monte Carlo simulation is used to sample across the statistical spread of household demand and produce probabilistic outcome distributions at the network level [8].

D. Study Region

This study focuses on Virginia's Eastern Shore as shown in Figure 1, a geographically isolated peninsula bounded by the Chesapeake Bay and the Atlantic Ocean. This isolation produces a constrained distribution network with limited interconnection redundancy, making the region sensitive to demand-side stress and a useful testbed for demand response evaluation. The area's rural population and demographic profile are key inputs to the behavioral adoption model developed in Section III-B.

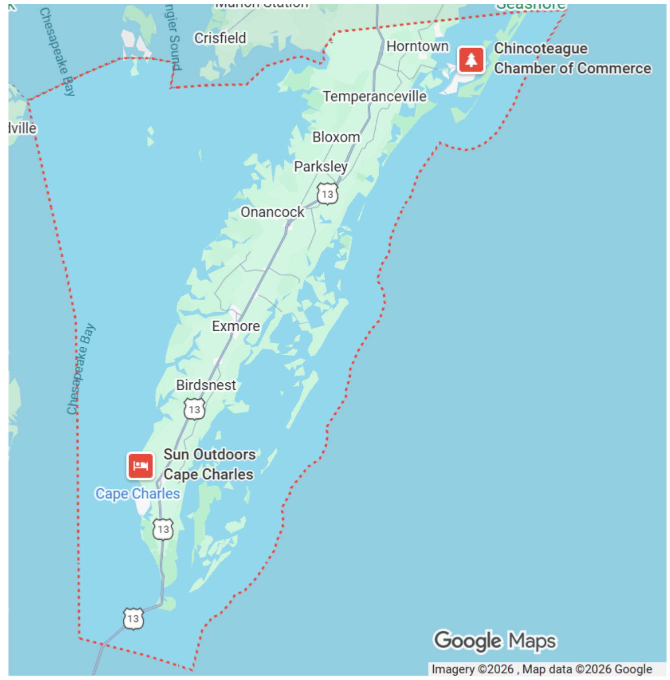

Figure 1. Study region: Eastern Shore, Virginia.

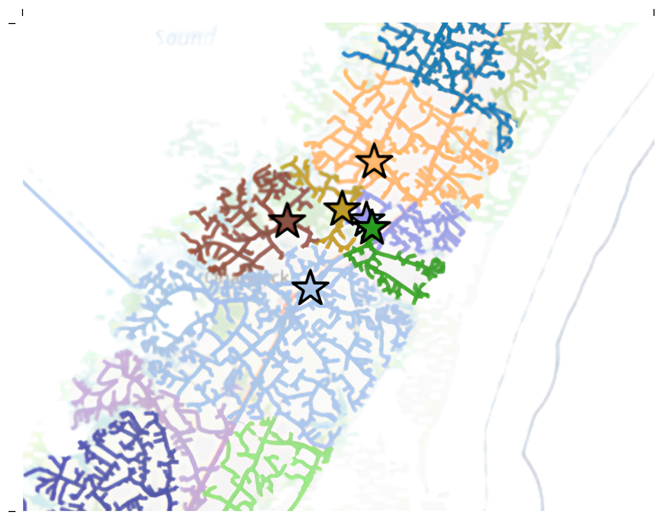

Figure 2. A portion of the distribution network of the ESVA digital twin. Stars indicate substations.

III. Methodology

A. Digital Twin Architecture

The digital twin couples three components: a synthetic household population, a set of hourly load profiles, and a synthetic distribution network. Gallagher et al. [9] provide the synthetic population for households across the United States. Then, Thorve et al. [10] construct hourly load profiles for each household, considering regional differences in climate, dwelling type, and demographic characteristics. Synthetic distribution network topology is derived from a published study by Meyur et al. [11] as shown in Figure 2. Together, these components produce a spatially and behaviorally explicit representation of residential electricity demand on the Eastern Shore under baseline and counterfactual policy scenarios, accounting for $28,221$ households assigned to 21 disjoint substation networks. Figure 3 diagrams the sequence of digital twin construction, PBDR experimentation, and analysis used in this study to extract findings.

B. Demographic-Based Stochastic Adoption Model

Before synthetic demand profiles are altered, empirical data are used to approximate which households change their

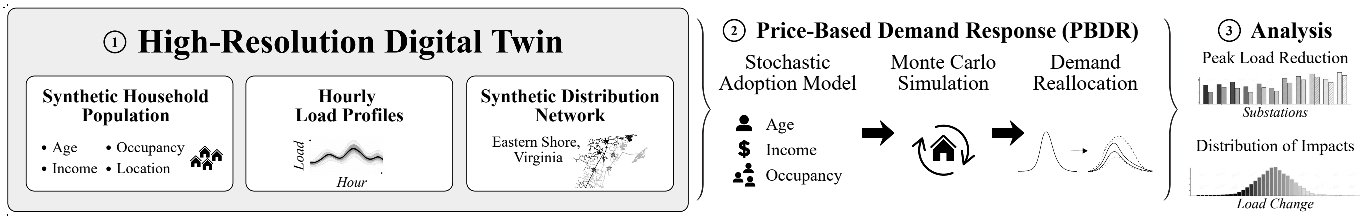

Figure 3. Overview of the methodology, integrating a high-resolution digital twin of Eastern Shore, Virginia, with a Price-Based Demand Response framework to estimate substation-level peak load reductions and load shift distributions.

behavior in response to demand response policies.

Household participation in a demand response policy is modeled as a binary outcome $D_i \in \{0, 1\}$, where $D_i = 1$ indicates participation. The probability of participation is specified as a logistic function of three demographic predictors: household maximum age a_i, annual income y_i, and household occupancy n_i. These factors are found by the qualitative analysis of Sridhar et al. [12] to consistently differentiate between households based on demand response enrollment. Marginal distributions for age group, income group, and household occupancy have been adapted from [13]–[16]. Table I shows the mapping of these characteristics to probability values β.

Table I
MARGINAL DISTRIBUTIONS FOR ADOPTION PROBABILITY BY DEMOGRAPHIC FACTOR.

Factor	Category	β
Maximum Age (years)	18-35	0.2
	36-65	0.5
	66+	0.1
Income ($/year)	0-49,999	0.2
	50,000-99,999	0.5
	100,000+	0.3
Number of Occupants	1-2	0.2
	3-4	0.4
	5+	0.5

To form a composite score in the absence of reliable empirical estimates of relative effect sizes, equal weights are assigned across the three factors following the 1/N heuristic:

$$\text{score}_i = \frac{\beta_{\text{age}} + \beta_{\text{inc}} + \beta_{\text{occ}}}{3} \tag{1}$$

The score is then mapped to a household-level adoption probability via a shifted logistic function,

$$P_i = \sigma(\text{logit}(\text{score}_i) + \alpha) \tag{2}$$

where α is an intercept chosen to calibrate the mean predicted probability across all households to a target adoption rate p_{overall}:

$$\frac{1}{N}\sum_{i=1}^{N} \sigma(\text{logit}(\text{score}_i) + \alpha) = p_{\text{overall}} \tag{3}$$

The target adoption rate p_{overall} is assigned according to policy design scenarios in Section III-C.

This approximation of the joint distribution constitutes a *reduced-form behavioral approximation* rather than a structural model of utility maximization. This approach is common in energy system modeling [17] and in policy simulation contexts where individual-level preference data are unavailable but aggregate response patterns are well-documented in the literature. This methodology is analogous to reduced-form diffusion frameworks such as the Bass model [18], which similarly parameterize population-level adoption without requiring individual-level preference data.

C. Policy Design Scenarios

Two demand response policy structures are evaluated, corresponding to the hour-based and day-based penalty mechanisms reviewed in Section II.

1) Time-of-Use (TOU) Pricing: The TOU scenario imposes time-differentiated volumetric electricity prices with a daily peak window and a complementary off-peak period. Peak hours are defined as 17:00–23:00 on weekdays. A similar peak window is used in the Belmont Light TOU pilot structure [3], but this study placed the peak ∼3 hours later in the day to account for the peaks observed in the load profiles. Households that participate in the TOU scenario on a given day shift 15% of load in the peak window to the 5 off-peak hours immediately before the peak window.

The behavioral adoption model for this policy treatment used $p_{\text{overall}} = 0.27$ [19].

2) Tempo-Style Red-White-Blue Day Pricing: The Tempo scenario is modeled on the French EDF Tempo tariff structure, in which days are classified into three tiers: Blue days (standard prices), White days (slightly above standard prices), and Red days (substantially elevated prices). The greatest load-shifting events are triggered only on Red days, limiting the behavioral burden on participating households relative to TOU. In contrast to the Tempo policies that assigned Red days during the winter, this study assigns the greatest demand shifts during the summer, where demand peaks are most likely to reach critical levels. On Red days, consumers' behavioral response to elevated prices is modeled as a 45% reduction in demand across all hours of the day, where displaced demand is uniformly reallocated to Blue days within the week. The same method is employed on White days, but a milder reduction of 15% is used. These reduction rates are supported by a 2009 study by Batlle and Rodilla [20].

The behavioral adoption model for this policy treatment used $p_{\text{overall}} = 0.09$ [21], [22].

3) Preservation of Total Daily Energy: Demand response is modeled as intertemporal load shifting rather than curtailment. For each participating household i on a TOU demand response day, the total daily energy consumption $E_i = \sum_{h=1}^{24} \ell_{i,h}$ is held constant before and after shifting, where $\ell_{i,h}$ denotes load in hour h. Load is redistributed from peak-designated hours to off-peak hours, preserving the household's total daily energy balance. The same conservation rule is applied to Tempo days using a larger scale. For each participating household i on a Tempo demand response day, the total weekly energy consumption $E_i = \sum_{d=1}^{7} \ell_{i,d}$ is held constant before and after shifting, where $\ell_{i,d}$ denotes load on day d.

D. Simulation Framework

1) Monte Carlo Runs: Each policy scenario is simulated over $M = 10$ independent Monte Carlo trials. In each trial, household participation is determined stochastically via Bernoulli draws from the calibrated logistic model, and the resulting aggregate load curve is computed for all days in the simulation period. With $N = 28{,}221$ household draws per trial, aggregate sampling variance is small. The observed cross-trial standard deviation of mean substation peak reduction ranges from 0.5 to 4.5 kWh against means of 7 to 116 kWh, yielding coefficients of variation between 4% and 12% at $M = 10$. Relative variability is highest at substations with low mean reductions, where the absolute standard deviation remains below 1 kWh.

2) Stability Diagnostics: Simulation stability is assessed by examining the convergence of key outcome statistics as a function of M. Specifically, the running mean and standard deviation of peak demand reduction and substation overload probability are tracked across trials, and M is considered sufficient when these statistics stabilize within a tolerance of $\pm 0.5\%$ of the final-trial estimate.

3) Outcome Metrics: Two primary outcome metrics are computed from the simulation output:

1) **Peak Demand Reduction.** The mean and distributional percentiles of the reduction in coincident peak demand at the substation level, relative to the no-policy baseline, expressed in kilowatts and as a percentage of baseline peak.
2) **Load Variance and Ramping.** The hourly variance of aggregate system load across the simulation period, computed both within-day and across-day, as well as maximum hourly change in demand. Demand response programs may flatten the peak but induce rebound peaks; this metric captures such redistribution effects.

IV. Results

A. Tempo

Tempo tariffs are simulated over one week in July. On the Red day, an average reduction in peak load of 4.1% is achieved across substations with a standard deviation of 0.79%. Across all days, the average peak load rose by 0.4%; because the absolute system peak coincides with the Red day (i.e. is chosen optimally) and is reduced across all 21 substations, this increase reflects demand displaced from Red days rather than a deterioration in system performance. This is evidenced by the 17.3% decrease in across-day variance compared to baseline. Within-day variance and maximum ramp rate did not change significantly. Participation rates range from 6.73-10.43% across substations, with a mean of 8.94%.

Figure 4 illustrates how the Tempo policy reduces peak demand on Red and White days while slightly increasing peaks on Blue days. The weekly load profile retains the shape of the baseline scenario. Additionally, optimal selection of Red, White, and Blue days is only feasible in hindsight.

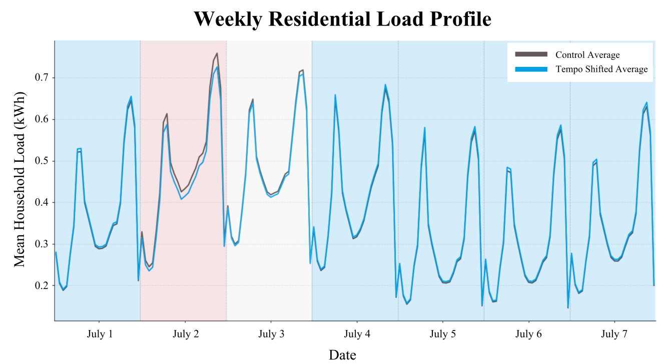

Figure 4. Residential load over a week using Tempo policy.

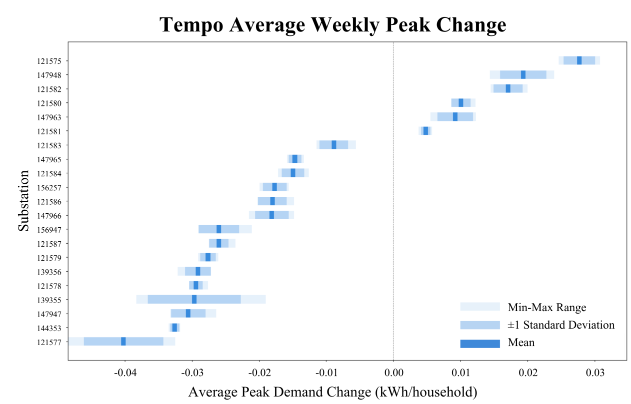

Figure 5. Change in per-household average demand peak over a week using Tempo policy.

Figure 5 visualizes the average change in household peak demand across the entire week, broken down by substation. The data are extracted from 10 Monte Carlo runs. After reallocating Red and White day demand, certain substations' average peak demand actually increases even though their absolute maximum decreases. Although this reduction in absolute peak load and ensuing reduction in across-day variance is a benefit to the system, these results reveal an important consideration: behavior at each substation may vary significantly,

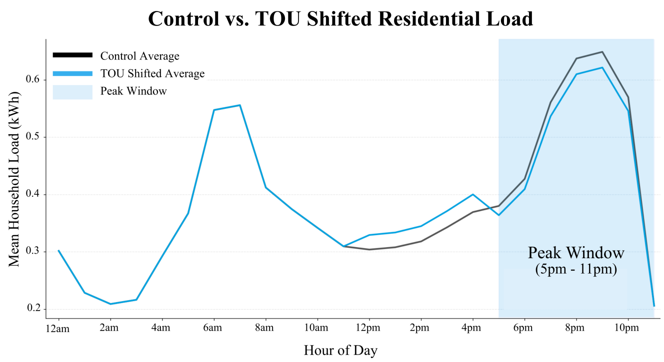

Figure 6. Shifted residential load in Time-of-Use policy.

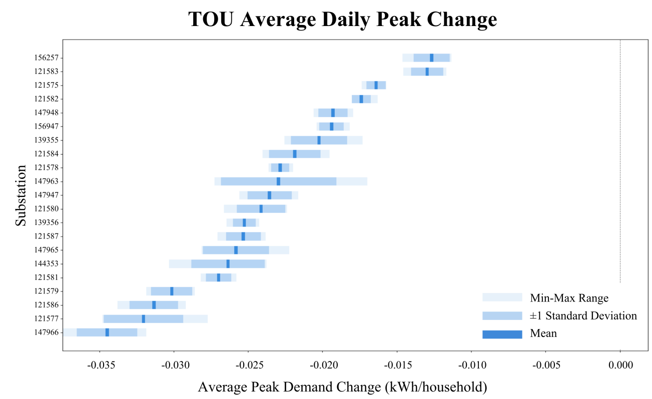

Figure 7. Change in per-household average daily demand peak using Time-of-Use policy.

and substation-level outcomes may be obscured if planning is conducted only at the system-wide level. When planning a dynamic pricing schedule, policymakers must protect the ongoing reliability of all individual substations while pursuing system-wide benefits. Substation failure, in the extreme case, is not an acceptable risk no matter the potential magnitude of system-wide load reduction.

B. Time-of-Use

The Time-of-Use tariff achieved a mean reduction of 3.6% of peak load at the substation level, ranging from 2.5-4.1%. Participation rates varied from 13-30% by substation, with a mean participation rate of 24.4%. TOU tariffs reduced the mean maximum ramp rate by 6.5% and reduced average within-day variance by 12.5%. Across-day variance did not change significantly. In Figure 6, demand is reduced during the peak window and increased in the hours before the peak window.

Figure 7 calculates the total amount of demand during the peak hour of each day for each substation. Then, this peak demand is averaged across days for each substation and divided by the number of households to normalize the results for substations serving differently sized residential groups. All substations experienced a decrease in peak demand under this policy, but some variation between substations still appears. While Tempo policies help mitigate the risk of infrequent, extreme demand events, TOU policies make demand curves less burdensome on the system every day.

V. Discussion

The quantitative results of this study are consistent with established findings in the PBDR literature. Both the TOU and Tempo scenarios produced peak load reductions of 4-5%, aligning with the 5-15% reduction documented in prior empirical studies of voluntary Time-of-Use programs [19]. The authors attribute the mild load reductions to the study region; higher income regions with higher consumption (Eastern Shore has a low rate of HVAC usage, for instance) tend to produce larger shifts. In this case study, participation rates did not vary to a statistically significant degree across substations. This result is attributable to the relative demographic homogeneity of Virginia's Eastern Shore, where income distributions and household compositions are comparatively uniform across substation assignments. In regions with greater socioeconomic diversity, particularly urban areas with sharp demographic gradients, the model would be expected to reveal meaningful spatial heterogeneity in adoption and, consequently, in substation-level load relief. The absence of such variation here should therefore be interpreted as a property of the study region rather than a limitation of the framework.

The novel contribution of this work is not in the magnitude or method for achieving these shifts, but in the resolution at which they are observed. Our framework derives system-level demand reductions from individual household decisions, linking adoption behavior, demographic characteristics, and load-shifting actions to measurable changes at specific substations. The fact that increased granularity does not reduce the accuracy or precision of the results validates the approach, demonstrating that well-established demand response patterns can be reproduced in a model operating at the household and substation levels. This reveals how energy consumption is spatially distributed across distribution infrastructure—information that is lost when residential demand is aggregated to the system scale.

For policymakers and utility planners, this capability addresses a persistent gap between policy decisions and infrastructure outcomes. Standard demand response evaluations treat load reductions as uniform across a service territory, but distribution networks are not uniform. Some substations serve dense clusters of high-consumption households while others operate well within capacity, and demographic groups with distinct responses to PBDR programs may cluster by substation. Our framework integrates physical infrastructure into policy evaluation and models behavioral responses at the correct unit of analysis: the household. This makes it possible to identify which substations are most exposed to demand-side stress and to direct limited capital and programmatic resources accordingly. A key limitation of this digital twin is its input data: the authors rely on synthetic data sets from 2014, which may not capture changes in population demographics and household electricity consumption. Our twin also focuses exclusively on residential consumption, which represents an increasingly smaller share of national electricity demand.

Several extensions merit further investigation. Adoption behavior, for instance, could be modeled in greater detail by incorporating heterogeneous response patterns, including variation in both the magnitude and timing of load shifting to better reflect differences in consumer sensitivity to price signals. Furthermore, simulating a longer time horizon permits analysis of "decision fatigue" and its impact on participation. Although Tempo may seemingly induce less load shifting than TOU, it places a lower cognitive load on consumers, which may be beneficial in the real-world. The framework could also be tested in urban distribution networks, where load density and demographic heterogeneity pose distinct planning challenges. Additionally, integrating the transmission network could enable analysis of system-wide capacity constraints, while incorporating distributed energy resources such as electric vehicles and rooftop solar would provide insight into the drivers of localized demand flexibility. Finally, future simulations could examine system performance under extreme conditions, such as prolonged cold events or heat waves, incorporating outage risk modeling to assess the extent to which demand response can mitigate service disruptions.

VI. Conclusion

This study developed a high-resolution digital twin to evaluate the effects of price-based demand response across Virginia's Eastern Shore. By modeling energy consumption at the household level and linking it to a synthetic distribution network, the framework captured localized load dynamics by substation that are not visible in more aggregated approaches. The results show that PBDR can reduce infrastructure stress without requiring immediate capacity expansion, and that established demand response effects can be reproduced at greater granularity without loss of fidelity. These findings highlight that grid resilience is not solely a property of physical infrastructure but also depends on consumer behavior, which policymakers can shape. The framework is designed to be portable to other regions and stress scenarios across the United States.

Acknowledgment

The team would like to thank all who have supported our research. This includes our advisor, Henning Mortveit, who supervised the project over the course of the 2025-2026 academic year. Additionally, our research extends the work of Swapna Thorve (Amazon) and Rounak Meyur (PNNL), both of whom offered their expertise and data for our research. We would also like to thank Bill Shobe (UVA), and the University of Virginia Environmental Institute. This work was sponsored in part by the NSF Grant "Focused CoPe HUB Coastal Futures: Building Capacity for Data-driven Adaption in Rural Coastal Communities" (Award #2053013).

References

[1] C. for Rural Affairs, "Load growth and the transmission system," Center for Rural Affairs, Tech. Rep., 2025.

[2] J. e. a. Rand, "Queued up: 2025 edition," Lawrence Berkeley National Laboratory, Tech. Rep., 2025.

[3] K. Bleau, "How load research informed belmont light's rate design," *American Public Power Association*, 2025.

[4] V. C. Pandey, N. Gupta, K. R. Niazi, A. Swarnkar, and R. A. Thokar, "An adaptive demand response framework using price elasticity model in distribution networks," *Electric Power Systems Research*, 2022.

[5] O. of Electricity Delivery and E. Reliability, "United states electricity industry primer," U.S. Department of Energy, Tech. Rep., 2015.

[6] U. D. of Energy, "Benefits of demand response in electricity markets and recommendations for achieving them," U.S. Department of Energy, Tech. Rep., 2006.

[7] R. de Transport d'Électricité (RTE), "Schedule of tempo-type supply offerings," 2026.

[8] D. H. O. McQueen, P. R. Hyland, and S. J. Watson, "Monte carlo simulation of residential electricity demand for forecasting maximum demand on distribution networks," *IEEE Transactions on Power Systems*, 2004.

[9] S. Gallagher *et al.*, "Spew: Synthetic populations and ecosystems of the world," *Journal of Computational and Graphical Statistics*, vol. 27, no. 4, pp. 773–784, 2018.

[10] S. Thorve *et al.*, "High resolution synthetic residential energy use profiles for the united states," *Sci Data 10, 76*, 2015.

[11] R. Meyur *et al.*, "Ensembles of realistic power distribution networks," *PNAS Vol. 119, No. 42*, 2022.

[12] A. Sridhar *et al.*, "Toward residential flexibility—consumer willingness to enroll household loads in demand response," *Applied Energy 342*, 2023.

[13] M. e. a. Nicolson, "Are consumers willing to switch to smart time of use electricity tariffs? the importance of loss-aversion and electric vehicle ownership," *Energy Research & Social Science, Volume 23*, 2017.

[14] P. Li, I. Keppo, M. Xenitidou, and et al., "Investigating uk consumers' heterogeneous engagement in demand-side response," *Energy Efficiency*, vol. 13, pp. 621–648, 2020.

[15] A. Faruqui and S. Sergici, "Household response to dynamic pricing of electricity: a survey of 15 experiments," *J Regul Econ 38, 193–225*, 2010.

[16] A. Faruqui *et al.*, "The impact of dynamic pricing on residential and small commercial and industrial usage: New experimental evidence from connecticut." *The Energy Journal*, vol. 35, 2014.

[17] A. Ball-Burack *et al.*, "Assessing the behavioral realism of energy system models in light of the consumer adoption literature," *Renewable and Sustainable Energy Reviews*, vol. 211, 2025.

[18] F. Bass, "A new product growth for model consumer durables," *Management Science*, vol. 15, pp. 215–227, 1969.

[19] M. Nicolson et al., "Consumer demand for time of use electricity tariffs: A systematized review of the empirical evidence," *Renewable and Sustainable Energy Reviews*, vol. 97, 2018.

[20] C. Batlle and P. Rodilla, "Electricity demand response tools: current status and outstanding issues," *Comillas University of Madrid, Spain*, 2009.

[21] B. Martucci, "Demand response programs improving, but customers remain wary: report," *Utility Dive*, 2025.

[22] B. Loeff, "Demand response: Are residential customers ready to participate?" *American Public Power Association*, 2024.