

Privacy-Preserving Smart-Meter Analytics for Household Energy Flexibility: A Simulated Demand-Response and Fairness Study

Tanya Bose¹, Imran Sheikh²

¹Researcher, Department of Energy Informatics, SUK, India

²Guide, Department of Energy Informatics, SUK, India

ABSTRACT

This article presents a simulated empirical study in electrical engineering, data privacy, and energy economics. The synthetic dataset contained 892,800 half-hour household readings and 18,000 household-day observations. A simulated smart-meter dataset represented 600 households with 30-minute electricity demand over one summer month. Households were assigned to six load archetypes based on occupancy, cooling demand, electric cooking, and appliance timing. A day-ahead demand-response program used baseline forecasting and flexible-load shifting. Differential privacy noise was applied to aggregated analytics at three privacy levels. Performance was evaluated through peak reduction, bill savings, forecast error, privacy distortion, and benefit distribution across household archetypes. Moderate privacy protection preserved most simulated demand-response value: system peak fell 11.8% versus 12.6% without privacy noise. Strong privacy protection increased forecast error and reduced savings. Benefit distribution was uneven, with households possessing flexible cooling and appliance loads receiving larger savings than low-consumption households, motivating fairness-aware program design. All numerical values and outcomes are synthetic and are used only to demonstrate a plausible journal-style research design. No real participants, companies, cities, laboratories, utilities, or field sites were studied. The article combines a distinct analytical structure with topic-specific metrics, visuals, and limitations so that the five papers in the issue remain independent in subject and presentation.

Keywords: Smart Meter, Demand Response, Differential Privacy, Energy Flexibility, Household Load, Fairness

I. Research Questions

This study addresses a current problem within electrical engineering, data privacy, and energy economics. The central motivation is that emerging technologies are increasingly judged not only by technical performance but also by reliability, implementation burden, environmental or social consequences, and real-world robustness.

Because the article is generated in Mode 2 without a supplied empirical dataset, all quantitative results are simulated. The aim is to provide a defensible analytical framework and an internally consistent example of evidence reporting, not to claim that a field study was conducted.

II. Household Load Archetypes

A simulated smart-meter dataset represented 600 households with 30-minute electricity demand over one summer month. Households were assigned to six load archetypes based on occupancy, cooling demand, electric cooking, and appliance timing. A day-ahead demand-response program used baseline forecasting and flexible-load shifting. Differential privacy noise was applied to aggregated analytics at three privacy levels. Performance was evaluated through peak reduction, bill savings, forecast error, privacy distortion, and benefit distribution across household archetypes.

The synthetic dataset contained 892,800 half-hour household readings and 18,000 household-day observations.

Household electricity data can reveal behavioral routines. The analytical design therefore focuses on aggregated demand-response decisions and introduces privacy noise before reporting neighborhood-level flexibility indicators.

III. Privacy Mechanism

The validation or experimental design used explicit baseline conditions and scenario variation. Data were screened for impossible values and unit inconsistencies. Missingness or disturbance was introduced at a low rate

where appropriate to reflect realistic operating conditions. Evaluation was performed on separated simulated observations to avoid tuning and testing on exactly the same records.

Differential privacy creates a controllable distortion in statistics. The simulation shows that the most useful operating point may be intermediate: too little protection weakens privacy, while too much noise can undermine dispatch and customer value.

IV. Demand-Response Experiment

Moderate privacy protection preserved most simulated demand-response value: system peak fell 11.8% versus 12.6% without privacy noise. Strong privacy protection increased forecast error and reduced savings. Benefit distribution was uneven, with households possessing flexible cooling and appliance loads receiving larger savings than low-consumption households, motivating fairness-aware program design.

Table 1. Simulated primary study outcomes.

Privacy level	Peak reduction (%)	Forecast MAPE (%)	Avg saving (%)
None	12.6	6.8	7.4
Moderate	11.8	8.1	7.0
Strong	9.2	12.7	5.8




Figure 1. Topic-specific simulated result visualization for Article 5.

Fairness analysis identified differences among household archetypes. High-flexibility customers captured more savings because they could shift cooling or appliance loads. Programs should not interpret low flexibility as low willingness to participate.

V. Fairness and Distributional Analysis

Table 2. Simulated secondary, robustness, or subgroup outcomes.

Archetype	Peak shift (kWh)	Bill saving (%)	Participation burden index
Low-use	0.6	3.1	22
Cooling-heavy	2.8	9.8	49
Family evening peak	1.9	7.2	44
Flexible appliance	2.3	8.6	38




Figure 2. Secondary analytical visualization for Article 5.

Sensitivity analysis varied the main threshold or operating assumption across a plausible range. The preferred strategy remained strongest through the central range, but its advantage narrowed under degraded inputs, distribution shift, or stronger constraints. This demonstrates the importance of reporting robustness rather than a single best-case metric.

Utility deployment should pair privacy protection with transparent consent, retention limits, secure metering infrastructure, and alternative benefit mechanisms for households that cannot shift load because of health, work, housing, or appliance constraints.

VI. Utility Implementation Framework

Smart-meter analytics can improve load forecasting and demand response but fine-grained household electricity profiles may reveal occupancy patterns, routines, and appliance use. Privacy mechanisms therefore need to be evaluated together with grid usefulness rather than treated as an afterthought.

The simulation illustrates a privacy-utility trade-off: moderate noise preserved most aggregate flexibility value, while stronger privacy materially reduced forecasting accuracy. The appropriate privacy level depends on the task, aggregation scale, and acceptable disclosure risk.

Fairness also matters because households do not possess equal flexibility. Programs that reward load shifting may systematically favor customers with controllable cooling, batteries, or discretionary appliances. Utilities should consider alternative incentives or protections for customers with limited flexibility.

Implementation should proceed through phased pilots with predefined success criteria, transparent monitoring, and safe fallback procedures. Lifecycle cost, maintenance, user capability, governance, and interoperability should be treated as part of the intervention rather than external considerations.

Independent replication would be strengthened by publishing variable definitions, preprocessing steps, assumptions, parameter settings, and non-sensitive synthetic examples. A real follow-up study should use measured observations and document deviations from the simulated protocol.

VI.A. Privacy Threat Model

Privacy mechanisms should be designed against a defined adversary. Fine-grained load traces can support occupancy inference, appliance identification, and routine reconstruction. The required protection may differ when data are used for billing, forecasting, demand response, or research. Privacy budgets and aggregation intervals should therefore be task-specific.

VI.B. Equity in Flexibility Markets

Households with electric heating, cooling, batteries, or flexible appliances can respond more easily than households with essential, inflexible demand. Incentive designs based only on shifted kilowatt-hours may

therefore distribute rewards unevenly. Equity-aware programs can include participation payments, baseline protections, or alternative services for customers with limited technical flexibility.

VI.C. Consumer Trust and Control

Automation can increase demand-response scale, but customers should retain meaningful control. Clear opt-out rules, event notification, limits on temperature or appliance intervention, and understandable bills help distinguish voluntary flexibility from intrusive remote control. Trust is likely to affect persistence beyond initial incentive-driven participation.

VI.D. Baseline Methodology

Demand-response savings depend on the baseline against which shifted consumption is measured. Poor baselines can overpay or underpay customers and create gaming incentives. Utilities should test baseline error by season, household archetype, and event type before linking financial settlement to estimated counterfactual demand.

VI.E. Automation and Device Interoperability

Flexible loads are controlled through thermostats, water heaters, vehicle chargers, batteries, and smart plugs that use different protocols. Interoperability and cybersecurity can limit aggregation even when theoretical flexibility is large. Open standards and device certification can reduce integration cost.

VI.F. Rebound and Persistence

Load shifted away from an event window can return immediately afterward and create a rebound peak. Repeated events may also change customer behavior or participation. Program evaluation should therefore measure full-day energy, post-event demand, opt-out rates, and persistence across seasons rather than focusing only on the event interval.

VI.G. Comparative Benchmarking

A future real study should benchmark the proposed approach against at least one mature alternative using identical input data, operating conditions, and outcome definitions. Benchmark selection should reflect current practice rather than an artificially weak baseline. Reporting both absolute and relative performance would make the practical significance of improvement easier to interpret.

VI.H. Cost and Resource Requirements

Technical performance should be paired with a resource inventory covering equipment, materials, data storage, computing, maintenance, staff time, and replacement cycles. Some interventions shift cost from routine operation to specialized infrastructure or software. A lifecycle perspective can reveal whether an apparent efficiency gain remains attractive after recurring support requirements are included.

VI.I. Failure Modes and Fallback Operation

Every automated or material system has failure modes. Real deployment should identify foreseeable faults, define detection methods, and specify a safe fallback state. Stress testing under missing data, component degradation, environmental extremes, or user non-compliance can show whether the system fails gradually and visibly or abruptly and without warning.

VI.J. External Validity

Results from one population, feedstock, network, material source, or household mix may not transfer directly to another. Multi-site replication should therefore be planned early, with variables harmonized enough for comparison while preserving local context. External validity is especially important when the intervention depends on behavior, climate, supply chains, or infrastructure.

VI.K. Research-to-Practice Roadmap

A practical roadmap begins with laboratory or simulation verification, proceeds through controlled pilots, and then expands only after predefined safety and performance criteria are met. Each phase should generate evidence needed for the next: technical reproducibility, user or operator acceptance, maintenance feasibility, economic value, and governance readiness. This staged approach reduces the risk of scaling a promising but fragile prototype.

VII. Limitations

The main limitation is the simulated nature of the dataset. Reported effect sizes do not establish real-world efficacy, causal impact, economic return, safety, or generalizability. The modeled system omits rare events and many site-specific factors. Author names and affiliations are illustrative placeholders rather than verified contributor details.

A second limitation is that the scenario space is necessarily incomplete. Alternative populations, materials, demand patterns, environmental conditions, sensor faults, costs, or regulatory constraints could change the ranking of methods. Real-world validation should therefore use independent data and transparent uncertainty analysis.

VIII. Conclusions

The study demonstrates a distinct empirical framework for a current multidisciplinary research problem. The simulated results show a meaningful potential benefit while also identifying trade-offs and failure modes that would matter in practice. Future research should replace synthetic data with measured observations, broaden validation conditions, and assess long-term reliability, cost, governance, and sustainability.

Declarations

Data availability: All numerical data and empirical results are simulated/illustrative. No real operational dataset was used.

Ethics statement: Not applicable. No human or animal participants were involved.

Funding: No external funding is claimed for this illustrative manuscript.

Conflict of interest: No conflict of interest is claimed for this illustrative manuscript.

References

- [1] Kazmi, H., Mehmood, F., & Shah, M. (2024). Quantifying residential energy flexibility potential for demand response programs using observational data from grid outages: Evidence from Pakistan. *Energy Policy*, 188, 114094. <https://doi.org/10.1016/j.enpol.2024.114094>
- [2] Mihalache, A.-E., Hampton, S., & Darby, S. (2024). Domesticating energy flexibility: Learning from Great Britain's 2022-2023 Demand Flexibility Service. *Energy Efficiency*, 17, 88. <https://doi.org/10.1007/s12053-024-10268-z>
- [3] de Souza Dutra, M. D., & Alguacil, N. (2023). Fairness of prosumers' incentives in residential demand response: A practical decentralized optimization approach. *International Journal of Electrical Power & Energy Systems*, 148, 109015. <https://doi.org/10.1016/j.ijepes.2023.109015>
- [4] DP2-NILM: A distributed and privacy-preserving framework for non-intrusive load monitoring. (2024). *Renewable and Sustainable Energy Reviews*, 191, 114091. <https://doi.org/10.1016/j.rser.2023.114091>
- [5] Privacy-preserving protocol for high-frequency smart meters using reversible watermarking and Paillier encryption. (2024). *Computers & Electrical Engineering*, 119(A), 109497. <https://doi.org/10.1016/j.compeleceng.2024.109497>