Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

The economic value of load flexibility and its viability as a substitute for grid expansion

Valentin Praun
Energieinstitut an der Johannes Kepler Universität Linz
Johannes Reichl
Energieinstitut an der Johannes Kepler Universität Linz
Katharina Rusch
rusch@energieinstitut-linz.at

Energieinstitut an der Johannes Kepler Universität Linz

Article

Keywords:

Posted Date: August 31st, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10651024/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

The economic value of load flexibility and its viability as a substitute for grid expansion

Valentin Praun[a], Johannes Reichl[a,b], Katharina Rusch[a,c,*]

[a] *Department of Energy Economics, Energieinstitut an der Johannes Kepler Universität Linz, Altenberger Straße 69, 4040, Linz, Austria*

[b] *LIT Future Energy Lab, Johannes Kepler Universität Linz, Altenberger Straße 69, 4040, Linz, Austria*

[c] *Department of Economics, Johannes Kepler Universität Linz, Altenberger Straße 69, 4040, Linz, Austria*

Abstract

Power grid investments drive energy transition costs as electrification and distributed energy resources increase demand for electricity distribution capacity, making flexibility an increasingly important pillar of cost-efficient decarbonization. Despite growing policy interest in incentivizing flexibility, determining economically justified incentive levels remains an unresolved challenge. We develop a framework that derives the maximum economically viable incentive budget for flexibility directly from avoided grid expansion costs while explicitly accounting for uncertainty in future load conditions. An illustrative application to an Austrian low-voltage distribution grid section yields economically justifiable levels below 1 euro cent per kW shifted under realistic prediction accuracy of critical load situations. Even under perfect prediction, incentive levels increase only to approximately 3 cents per kW shifted, placing them at the lower end of the incentive levels reported as effective in the flexibility literature. More generally, the framework shifts the evaluation of flexibility from shifted energy volumes to avoided peak capacity, aligning it directly with the economic rationale underlying grid expansion decisions.

*Corresponding author

Email address: rusch@energieinstitut-linz.at (Katharina Rusch)

Contents

1 Introduction **3**

2 Main **5**

2.1 Framework 5

2.2 Application to a representative distribution grid setting 8

2.3 Scenario-based load trajectory 8

2.4 Threshold definition for demand peaks and oversupply 8

2.5 Forecasting 9

2.6 Potential avoidance of required grid capacity 10

2.7 Grid capacity expansion costs 12

2.8 Economically viable incentives for flexibility 12

3 Conclusion **13**

4 Methods **14**

4.1 Scenario-based load trajectories 14

4.2 Forecasting 15

4.3 Annualized investment budget 17

Appendix **23**

1. Introduction

Electricity grids are emerging as one of the central cost drivers of the energy transition. The electrification of transport, heating, and industry, combined with the expansion of decentralized renewable generation, is placing unprecedented pressure on distribution grids (Intergovernmental Panel on Climate Change, 2022; International Energy Agency, 2023; European Commission, 2020). Meeting these additional capacity requirements will require substantial investments in network infrastructure across Europe, e.g. Hoffmann et al. (2025) estimate aggregated costs in the EU until 2050 of 2.3 trn €.

Peak capacity (in kW or GW), rather than annual electricity volumes (in kWh or GWh), drives grid expansion costs in the transition to carbon neutrality (International Energy Agency, 2023; Brown et al., 2018). Distribution grids are therefore dimensioned to accommodate expected peak loads together with safety margins specified by technical planning standards (e.g. Austria's Technical and Organisational Rules (TORs) E-Control Austria, 2024). Because such peak events occur only infrequently, network assets remain substantially underutilized during most of their lifetime. Figure 1 illustrates this using the Load Duration Curve (LDC) of an exemplary Austrian village based on data from 1,292 households, showing that required grid capacity is determined by only a small fraction of periods with exceptionally high electricity demand.

In addition to conventional demand peaks, increasing deployment of distributed renewable generation, particularly Photovoltaic (PV) systems, also gives rise to negative peak loads. These occur when local generation exceeds electricity demand, causing reverse power flows into the higher-voltage grid. Distribution grids, originally designed for unidirectional power flows, may then experience transformer overloading, voltage rises and other operational challenges. Mitigating these effects, once again, requires substantial grid investments.

A large body of literature and policy discourse points to flexibility as a promising alternative to costly grid expansion. By shifting electricity consumption during critical periods, flexibility measures such as load-shifting can mitigate peak loads and thereby defer or avoid infrastructure investments (Palensky and Dietrich, 2011; Albadi and El-Saadany, 2007). As a result, flexibility is widely regarded as a key building block of a cost-efficient energy transition (European Commission, 2019).

This paper considers flexibility induced through economic incentives, which encourage consumers to adjust their electricity demand during periods of critical system stress in order to smooth the LDC. (Hernandez Arias et al., 2020; Hirth et al., 2022; Burger and Luke, 2017; Agency for the Cooperation of Energy Regulators (ACER), 2023; European Commission, 2024; Labandeira et al., 2017; Hofmann et al., 2025). By rewarding consumption shifts away from congested periods, incentive-based approaches preserve consumer choice while promoting a more efficient allocation of scarce grid capacity. Appendix A discusses incentive-based flexibility relative to technological or quantity-based peak mitigation measures.

In principle, economically efficient flexibility incentives would emerge endogenously from price signals reflecting the temporal scarcity of grid capacity, consistent with the logic of peak-load pricing (Steiner, 1957; Joskow, 2007). However, this mechanism does not materialize in electricity distribution systems. Distribution grids are regulated natural monopolies, and local grid capacity is not traded in competitive markets. Consequently, scarcity signals must be approximated through regulated tariffs or explicit monetary incentives. In practice, this requires a central authority, such as the national energy regulator, to determine economically justified incentive levels for shifting

electricity consumption.

Against this background, a fundamental question arises: what is the maximum economically viable cost of flexibility before grid expansion becomes the cheaper option? This threshold defines the upper bound of economically justifiable incentives for flexibility. If the incentives required to induce sufficient flexibility exceed the avoided grid expansion costs, flexibility measures cease to be a cost-effective substitute for infrastructure investment; if the incentives remain below this threshold, such measures can reduce grid costs. Appendix B provides the formal welfare-economic rationale. Importantly, the upper bound of economically viable incentives does not by itself determine the flexibility realised in practice. Realised demand adjustments depend on response demand elasticities, irrespective of whether responses are automated, market-mediated or decision-based. Establishing this upper bound is therefore a necessary first step before incorporating behavioural response constraints.

While previous studies have primarily evaluated flexibility either in terms of shifted energy volumes and consumer demand response (e.g., Faruqui and Sergici, 2010; Labandeira et al., 2017; Hirth et al., 2022) or from the perspective of welfare effects and efficient pricing (e.g., Steiner, 1957; Borenstein, 2005), grid expansion decisions are driven by peak capacity requirements rather than annual electricity volumes (e.g., Joskow, 2007; International Energy Agency, 2023). Consequently, the economically justifiable level of flexibility incentives cannot be inferred directly from existing approaches. Bridging this gap requires evaluating the grid value of flexibility in capacity rather than energy terms and directly translating avoided peak capacity into avoided grid expansion costs.

Accordingly, the objective of this paper is to develop a framework that quantifies the economic value of flexibility in electricity distribution systems by deriving the maximum economically justifiable incentive budget from avoided grid expansion costs. Central to the framework is the explicit consideration of uncertainty, whereby residual peak events determine the grid capacity that must still be installed despite demand-side interventions. The resulting framework provides a system-level benchmark for evaluating flexibility independently of assumptions about consumer response behaviour.

We illustrate the framework using an exemplary residential distribution grid section under different electrification scenarios. While the empirical application focuses on a representative low-voltage transformer, the framework is generally applicable across voltage levels and customer structures. Our results show that the economic value of flexibility is fundamentally bounded by the residual peak capacity: even substantial flexibility cannot eliminate the need for grid expansion if critical peak events remain unmitigated. These findings demonstrate that flexibility should be evaluated from the same capacity-based perspective that ultimately determines grid investment decisions.

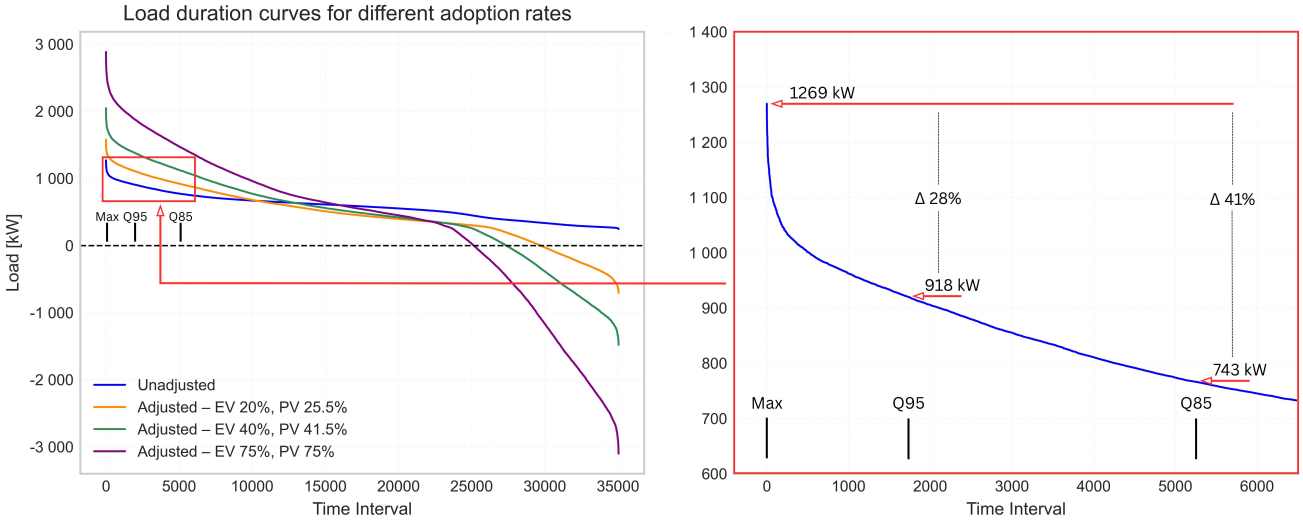


Figure 1 - Comparison of LDCs under different future Electric Vehicles (EV) and PV adoption scenarios: The LDC displays the aggregated grid load in kW for all 35,040 quarter-hour intervals of one full year, sorted in descending order. The blue curve represents the baseline load profile, while the adjusted curves include simulated additional EV consumption and PV generation for increasing adoption rates. Higher EV adoption increases electricity demand during charging periods and therefore raises the upper part of the LDC, whereas higher PV adoption reduces the grid load during periods of solar generation and can lead to negative load values. Consequently, increasing EV and PV adoption intensifies both positive and negative peak loads. The zoomed-in section on the right highlights the upper part of the baseline LDC. Within the first 1,752 quarter-hour intervals, corresponding to 5% of all intervals in the year, the load decreases from 1,269 kW to 918 kW, representing a reduction of about 28%. Within the first 5,256 quarter-hour intervals, corresponding to 15% of the annual intervals, the load further decreases to 743 kW, equivalent to a reduction of about 41%. Such intervals are particularly relevant from a grid-planning perspective, as they represent the range in which a comparatively small amount of flexibility can reduce the required peak grid capacity most effectively.

2. Main

2.1. Framework

The economic value of flexibility arises from its potential to avoid costly peak-driven capacity expansion. While network tariffs are typically defined for broad groups of consumers, the capacity constraints motivating flexibility are inherently local. Effective interventions should therefore target the specific grid section where these constraints arise.

The proposed framework accommodates this spatial heterogeneity and can be applied to any grid section for which a common set of incentives is to be defined. In the empirical application below, we illustrate it using a low-voltage grid section. Figure 2 summarizes the framework.

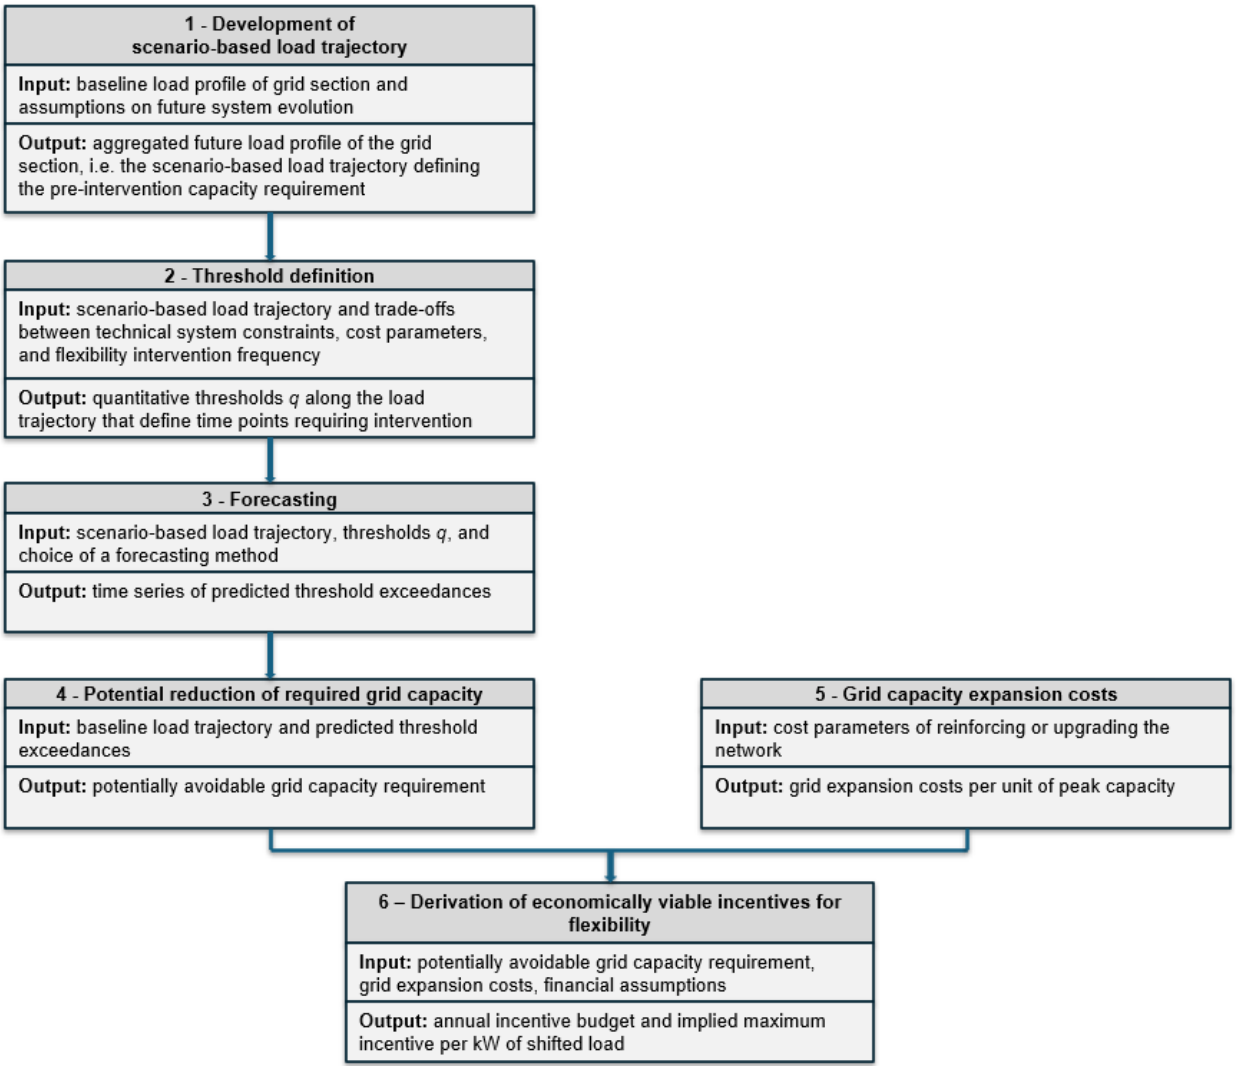


Figure 2 - Schematic representation of the framework for quantifying the economic value of flexibility: The approach links baseline load trajectories derived from load data and assumptions on future system evolution to threshold definition and forecasting of exceedance events. Forecasting performance determines residual peak capacity requirements, which are subsequently translated into avoided grid expansion costs and economically viable incentive levels.

First, observed high-resolution load data are aggregated at the target grid section to define the baseline load profile. This profile is then adjusted to reflect future electrification and distributed generation, yielding a scenario-based load trajectory that serves as the reference for evaluating flexibility measures.

Second, a threshold q defining critical load situations is specified for each scenario-based load trajectory. Determined by technical and economic considerations, the threshold identifies the peak events requiring intervention.

Third, the expected magnitude of threshold exceedances must be estimated before they occur. Because intervention decisions are made under uncertainty, they necessarily rely on forecasts, making forecasting performance a key determinant of the achievable reduction in required grid capacity. Figure 3 illustrates the resulting prediction errors. To quantify their economic relevance, in the empirical application (Section 2.8) we compare the empirical results with a *perfect-fit* benchmark

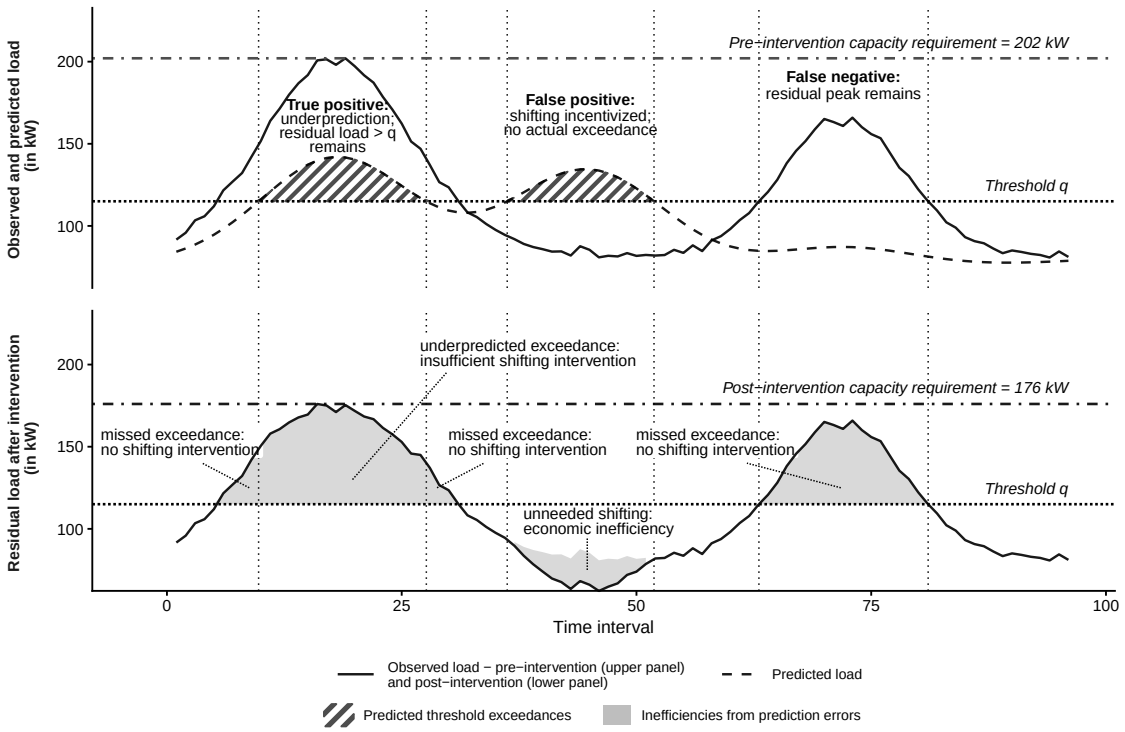


Figure 3 - Illustration based on synthetic example data of how different types of forecasting errors affect residual grid capacity requirements under forecast-based flexibility: The upper panel shows observed load (solid line), predicted load (dashed line), and a predefined intervention threshold q. Periods where the forecast predicts threshold exceedances trigger flexibility incentives (hatched areas). The Figure illustrates three types of prediction outcomes: correctly predicted exceedances, false positives (predicted exceedances without actual system stress), and false negatives (missed exceedances). Correctly predicted exceedances reduce residual peak demand, whereas false positives induce unnecessary interventions and false negatives leave peak events unmitigated. The lower panel shows the resulting residual load after intervention under the assumption that all correctly predicted exceedances are fully shifted. Residual peaks caused by missed exceedances or underprediction determine the post-intervention capacity requirement.

assuming all threshold exceedances are predicted correctly.

Fourth, predicted threshold exceedances are translated into avoided peak capacity. Assuming all correctly predicted exceedances are fully mitigated, the resulting capacity reduction represents an upper bound on the system-level value of flexibility, independent of behavioural responses or demand elasticities.

Fifth, reductions in peak capacity requirements are monetized using cost estimates for grid expansion, expressed in capacity-based infrastructure costs (e.g., €/kW). This step requires explicit assumptions about the cost of reinforcing or upgrading the grid at the relevant grid level.

Finally, avoided capacity additions are translated into avoided grid investment costs. Incentive-based flexibility is economically viable whenever the required incentive payments remain below these avoided costs, providing a direct decision criterion for evaluating flexibility as an alternative to grid expansion.

2.2. Application to a representative distribution grid setting

To illustrate the framework using empirical data, we consider an application based on the Austrian electricity system. Recent estimates indicate that accommodating electrification and renewable integration will require an additional 40–45 bn € of distribution grid investment by 2040 (Schwalbe et al., 2024; Oesterreichs Energie, 2025), alongside approximately 9 bn € for transmission expansion (Austrian Power Grid (APG), 2025). Together, these investments are expected to substantially increase grid costs, potentially doubling network tariffs or leading to sustained annual increases of around 3–5% in real terms (see e.g. Azarova et al., 2018, for an explanation of how investment costs enter network tariffs).

Against this background, flexibility is increasingly discussed as a policy instrument in the Austrian regulatory debate. The empirical application focuses on a low-voltage distribution grid section at the level of a local transformer, representing a typical aggregation of residential demand. This level constitutes a natural unit of analysis because local transformers often represent critical bottlenecks in distribution grids. At the same time, the cost of grid expansion at this level is typically high relative to the associated load, making it a particularly promising target for flexibility interventions.

2.3. Scenario-based load trajectory

The empirical application is based on aggregated electricity consumption from 1,292 Upper Austrian households, representing an exemplary low-voltage distribution grid supplied by a local transformer. Load data were recorded at 15-minute intervals from July 2017 to September 2018. Household characteristics were used for data imputation where required, as is detailed in Appendix C.

We construct three scenario-based load trajectories representing increasing adoption of rooftop PV systems and EVs. The scenarios assume adoption rates of 25.5%/ 20%, 41.5%/ 40%, and 75%/ 75% for PV and EVs, respectively, covering plausible medium- to long-term electrification pathways (Schwalbe et al., 2024). For each scenario, corresponding PV generation and EV charging profiles are incorporated into the baseline load to obtain adjusted demand trajectories. Details of the underlying assumptions and data construction are provided in Methods (Sections 4.1 and Appendix D).

The resulting scenario-based load trajectories are shown in Figure 1. Compared with the baseline profile, increasing EV adoption raises positive peak demand, whereas increasing PV penetration extends the lower tail of the distribution through periods of local generation exceeding consumption. Consequently, the load distribution widens substantially, increasing both demand- and generation-driven capacity requirements.

2.4. Threshold definition for demand peaks and oversupply

Thresholds are defined for each scenario-based load trajectory to reduce peak capacity requirements while restricting interventions to rare, system-critical conditions that drive grid expansion costs. Consequently, a purely physical definition of peaks is not appropriate; in practice, threshold selection reflects both technical and economic considerations.

For the empirical application, thresholds are defined from the statistical characteristics of each scenario as the 85^{th} and 95^{th} quantiles of all positive loads. For the lower tail, the threshold is set to zero, reflecting the perspective of low-voltage distribution system operation, where local PV generation exceeding demand may induce reverse power flows and associated operational challenges

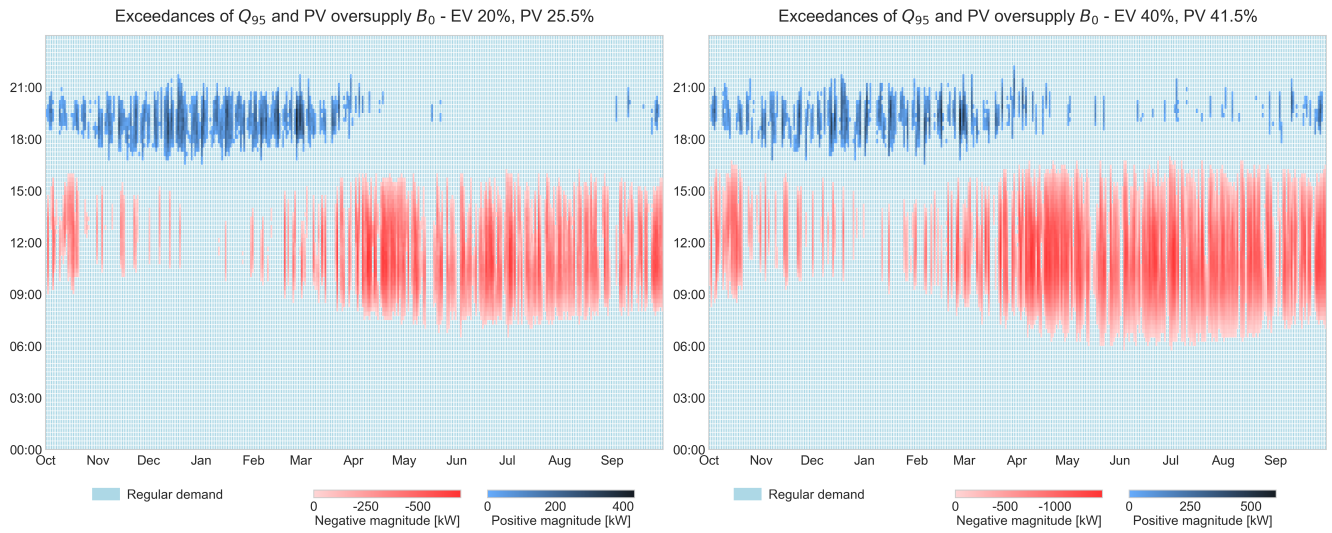


Figure 4 - Temporal distribution of threshold exceedances/ oversupply of adjusted electricity load for (20% EV, 25.5% PV) and (40% EV, 41.5% PV) adoption scenarios: The x-axis of the figure shows each day of the year, while the y-axis shows the 15-minute intervals within a day. Intervals where the simulated load exceeds the Q_{95} threshold are shown in dark blue, and intervals with PV oversupply where net consumption becomes negative are shown in red. The left panel shows the scenario of 20% adoption for EV and 25.5% for PV while the right shows an adoption of 40% EV and 41.5% PV. In both cases, exceedances cluster in the winter months. However, under the 40% EV and 41.5% PV scenario, their distribution becomes increasingly uniform throughout the year, indicating that the strong presence of EV whose charging requirements are largely constant throughout the year emerges as the primary driver of peak events. Consequently, these peaks are less influenced by seasonal factors such as increased heating demand in winter. For times of PV oversupply, midday hours are most strongly affected, as under the 40% EV and 41.5% PV scenario nearly all midday intervals experience negative net consumption due to high PV production while also under the 20% EV and 25.5% PV adoption scenario, negative consumption occurs on most days during summer as well as on selected days in winter during noon.

(Turitsyn et al., 2011; Lopes et al., 2007). Although this represents a stricter criterion than typically applied in practice, it provides a transparent benchmark for the illustrative application.

Accordingly, the framework considers both excessive demand and local generation exceeding demand as conditions requiring intervention. Figure 4 illustrates the temporal distribution of threshold exceedances for two representative scenarios.

2.5. Forecasting

To identify periods requiring flexibility interventions, we generate short-term forecasts of aggregated transformer load. Forecasts are produced for a 24-hour day-ahead horizon at 15-minute resolution and provide both continuous load predictions for estimating potential shifting volumes and probabilities of exceeding the upper thresholds or falling below zero. Within the proposed framework, forecasting serves as an input to the economic decision problem by identifying future grid-critical periods. Details of the forecasting methodology are provided in Section 4.2.

Figure 5 illustrates forecast outcomes for the scenario with 20% EV and 25.5% PV adoption. The left panel highlights false negatives, i.e. threshold exceedances that were not identified by the forecasting model. Missed exceedances above the Q_{95} threshold are shown in dark blue and missed negative-load events in red, with colour intensity indicating the magnitude of the missed event. These missed events determine the residual peak capacity that must still be accommodated by the

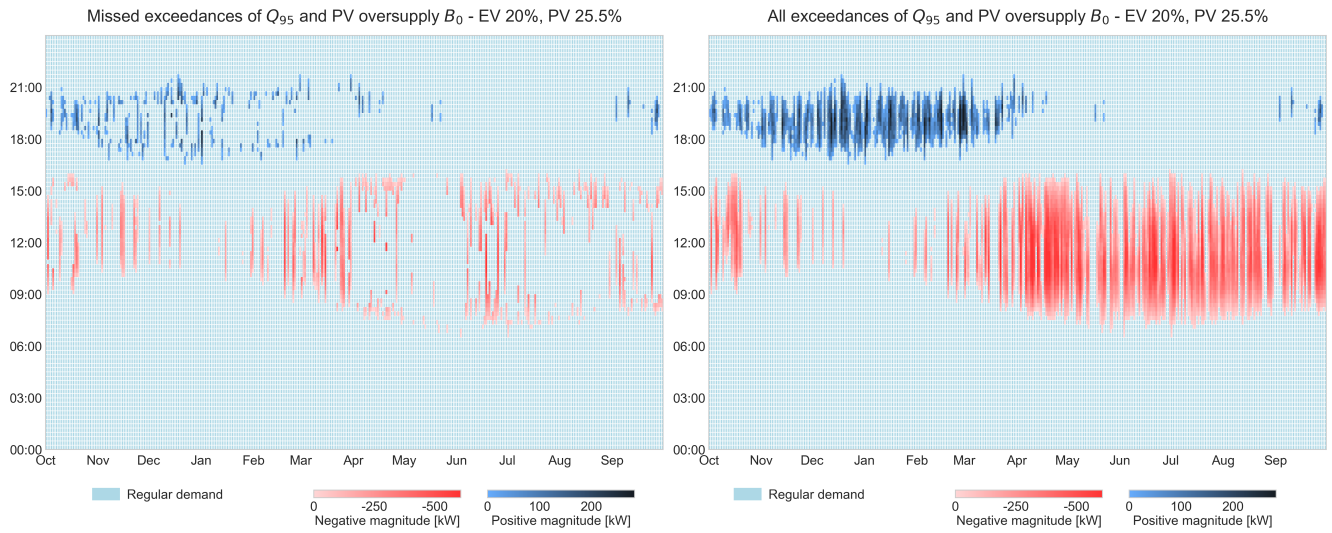


Figure 5 - Classification results false negatives only (EV 20%, PV 25.5%) under the Q_{95} threshold: The left panel shows all time intervals that the model incorrectly classified as non-events (false negatives) as derived from the continuous forecast: dark blue marks missed exceedances above Q_{95}, while red marks missed periods of PV oversupply. The right panel shows the corresponding true exceedances for comparison. To maintain the focus on the more critical false negatives for now, any false positives are shaded in light blue as if they were "normal" non-events. The false positive frequency is later shown in Figure 6. Overall, the model captures the majority of actual exceedances and oversupply periods. Most missed events which are left occur near the boundaries of exceedance hours for example, when the model predicts the event to begin slightly too late or end slightly too early. For the modelling accuracy of the continuous load forecast, rather than the derived binary classification, the model achieves an R^2 of 0.84.

grid despite flexibility interventions. The right panel shows the corresponding observed threshold exceedances for comparison. Results for the Q_{85} threshold are reported in Appendix E.

2.6. Potential avoidance of required grid capacity

We translate the forecasting results into system-level implications by recomputing the scenario-based load trajectories under the assumption that all correctly predicted threshold exceedances are successfully shifted. The remaining post-intervention capacity requirement therefore arises solely from missed exceedances or underpredictions, quantifying the avoidable peak capacity.

Figure 6 compares the pre- and post-intervention capacity requirements under the three electrification scenarios. Across all scenarios, forecast-based flexibility substantially reduces the required positive capacity, although only a comparatively small fraction of peak events is targeted. Under the Q_{85} threshold, positive capacity reductions range from 14.8% to 19.6%, whereas under the more selective Q_{95} threshold they range from 2.3% to 8.9%. Reductions in negative capacity remain smaller (9.0–15.6%), reflecting the greater difficulty of reliably identifying severe oversupply events. Across all scenarios, false positive frequencies remain low, ranging from 0.6–2.4% for positive peaks and 2.0–3.7% for negative-load events. Additional results are provided in Appendix F.

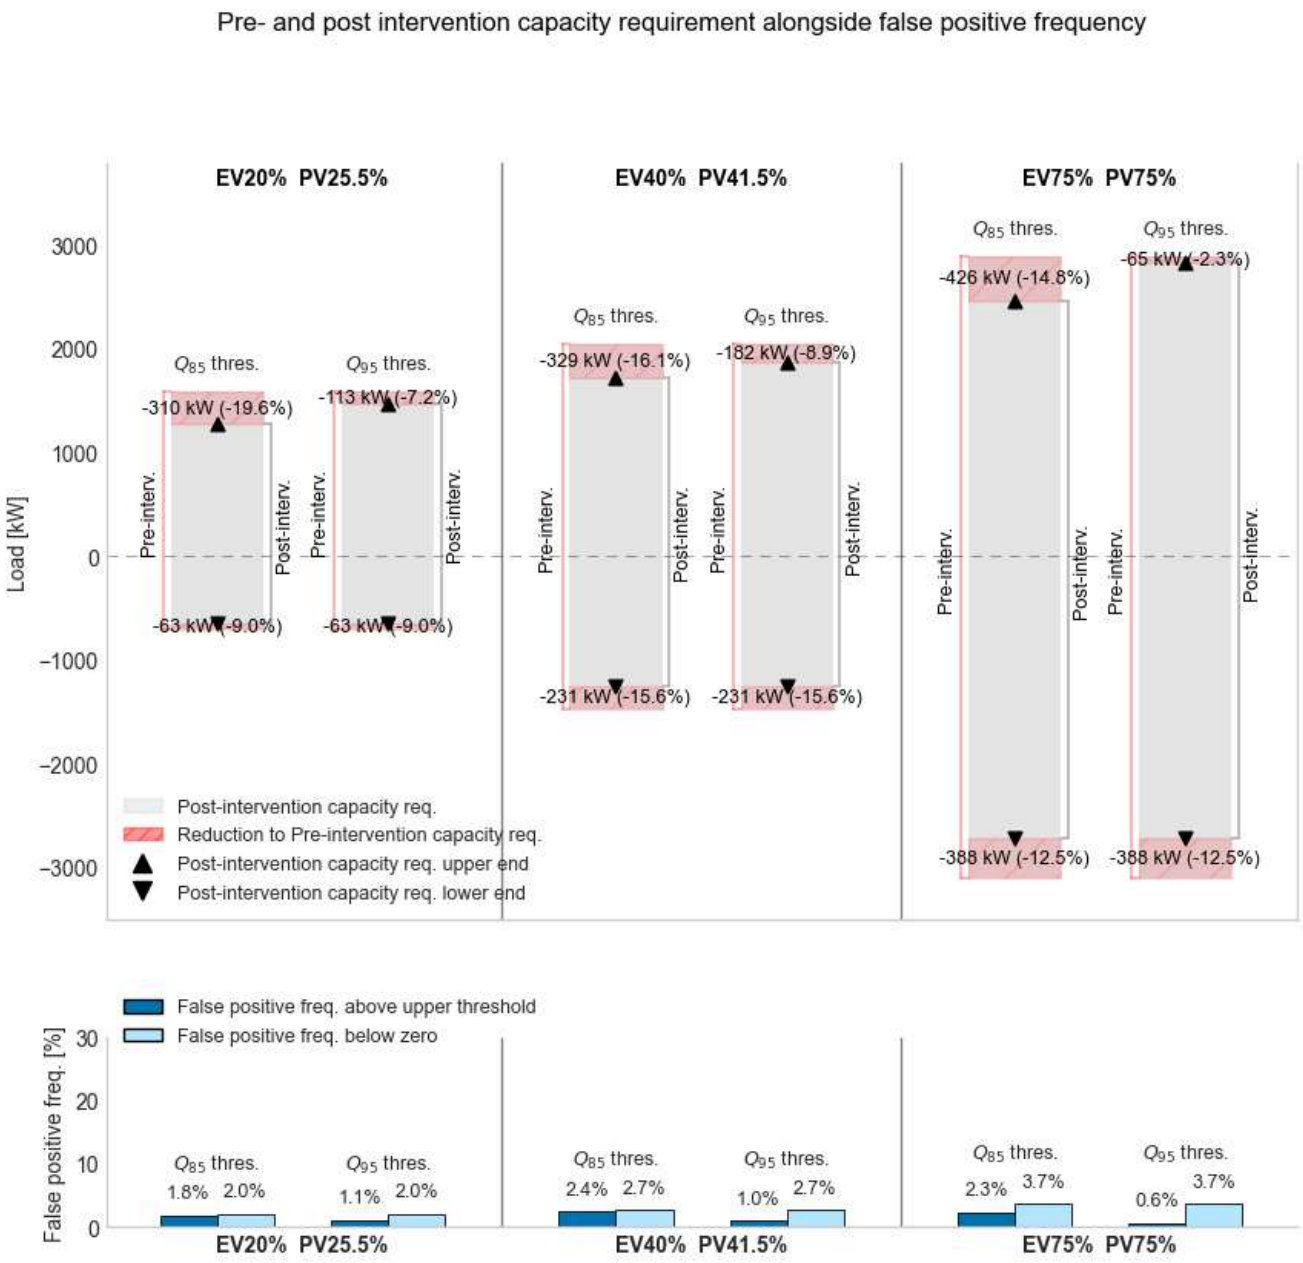


Figure 6 - Post-intervention capacity requirements under (20% EV, 25.5% PV), (40% EV, 41.5% PV), (75% EV, 75% PV) adoption scenarios: The grey bars show the post-intervention capacity requirement after applying the model. This post-intervention requirement is denoted by K^{new} as described in more detail in Appendix G. The red shaded areas illustrate the reduction in capacity relative to the pre-intervention capacity requirement as the baseline. For the lowest adoption scenario under the Q_{85} threshold, the required positive range is reduced by 19.6%, and the required negative range by 9.0%. For the higher Q_{95} threshold, the required positive range is reduced by 7.2% with the required negative capacity range staying the same as the target remains the same. For the remaining simulated scenarios the reduction in the magnitude of the post-intervention capacity requirement at the upper end remains comparable, though slightly decreasing. Under negative-load conditions, the potential for load curtailment at the lower bound remains comparatively limited and does not increase monotonically with rising adoption rates. It rises from 9.0% in the (20% EV, 25.5% PV) configuration to 15.6% in the (40% EV, 41.5% PV) scenario, before decreasing to 12.5% under the (75% EV, 75% PV) scenario. For completeness, the lower panel shows the false positive frequency for each target. For all targets, the frequency stays at minimal levels, remaining between 2.0% and 3.7% for negative load conditions, between 1.8% and 2.4% for exceedances above the Q_{85} threshold and between 0.6% and 1.1% for exceedances above the Q_{95} threshold.

2.7. Grid capacity expansion costs

Ideally, the economic value of flexibility would be quantified using location-specific grid expansion costs. Such assessments are routinely conducted by grid operators when evaluating local reinforcement decisions. However, such granular cost data are not available for the exemplary application considered here.

We therefore rely on the grid expansion cost estimates reported by Schwalbe et al. (2024), which quantify investment requirements for Austrian distribution grids under different electrification scenarios. Because our empirical application represents a low-voltage distribution grid, we use the corresponding low-voltage cost estimates.

The study estimates low-voltage grid investment requirements of 6.5 bn € under the 2030 pathway and 9.2 bn € under the 2040 pathway. Combining these estimates with the corresponding effective peak loads yields unit grid expansion costs expressed in €/kW, which form the basis of the subsequent economic valuation. Although heat pumps are not explicitly included in our simulations, whereas Schwalbe et al. (2024) do, expressing expansion costs per unit of peak capacity makes the analysis independent of the underlying composition of demand sources.

2.8. Economically viable incentives for flexibility

To derive economically viable incentive levels, we first determine the post-intervention capacity requirement K^{new}. Because grid expansion is driven by the maximum capacity that must be accommodated rather than individual exceedance events, the difference between the pre- and post-intervention capacity requirements represents the avoidable grid capacity (see Appendix G).

Using unit grid expansion costs (€/kW) from Schwalbe et al. (2024), avoided capacity is translated into avoided investment costs. These cost savings define the maximum incentive budget, which is annualized over a 20-year depreciation period and allocated across all forecast-triggered flexibility interventions, irrespective of whether they subsequently correspond to true or false positives. Details of the calculation are provided in Appendix G.

Table 1 reports the resulting incentive levels under the forecast-based post-intervention capacity requirement. Across all scenarios, annual incentive budgets remain modest relative to the underlying grid investment needs. Although the lower Q_{85} threshold targets a larger fraction of peak events than Q_{95}, the implied incentives per shifted kW differ only slightly because the larger budget is distributed across more flexibility requests.

Table 2 complements these results with a theoretical upper bound for flexibility incentives assuming all threshold exceedances are *perfectly* predicted and successfully shifted. In this case, the post-intervention capacity requirement is reduced to the selected thresholds, representing the maximum grid expansion cost savings achievable through flexibility. As expected, the perfect-fit benchmark yields substantially higher annual incentive budgets than the forecast-based configuration because no peak events are missed. Nevertheless, even under perfect information, the implied incentive per shifted kW remains modest because the amount of capacity that must be shifted increases sharply. The comparison demonstrates the economic relevance of uncertainty: missed or underpredicted peak events directly reduce the value of flexibility by increasing the residual capacity that must still be installed. In this framework, uncertainty therefore has a direct economic interpretation, as every missed peak event increases the capacity requirement that cannot be avoided through flexibility and must instead be provided through conventional grid expansion. This parallels the

Table 1 Simulated incentive levels considering post-intervention capacity results.

		2030 path	**2040 path**	**High scenario**
Q_{85}	Annual incentive budget (€/a)	6,881	16,112	17,663
	Incentive per household (€/a)	5.3	12.4	13.7
	Incentive per shifted kW (€)	0.007	0.005	0.002
Q_{95}	Annual incentive budget (€/a)	14,526	21,807	31,704
	Incentive per household (€/a)	11.2	16.9	24.5
	Incentive per shifted kW (€)	0.010	0.006	0.003

Table 2 Simulated incentive levels considering the perfect fit (all exceedances perfectly predicted in advance).

		2030 path	**2040 path**	**High scenario**
Q_{95}	Annual incentive budget (€/a)	44,283	80,993	154,891
	Incentive per household (€/a)	34.2	62.6	119.88
	Incentive per shifted kW (€)	0.029	0.017	0.012
Q_{85}	Annual incentive budget (€/a)	51,928	90,609	169,440
	Incentive per household (€/a)	40.2	70.1	131.1
	Incentive per shifted kW (€)	0.026	0.017	0.013

findings of Ruhnau et al. (2020), who show that the economic value of renewable generation forecasts depends not only on forecast accuracy but on the economic consequences of forecast errors.

3. Conclusion

This paper develops a framework for quantifying the economic value of flexibility by deriving the maximum economically justifiable incentive budget from avoided grid expansion costs. Central to the approach is the explicit treatment of uncertainty: residual peak events that remain unmitigated despite flexibility measures determine the grid capacity that must still be installed. The framework thereby links uncertainty, residual peak capacity, avoided infrastructure investment and economically viable incentive levels.

Applying the framework to an exemplary Austrian low-voltage distribution grid illustrates these relationships under plausible future electrification scenarios. Forecast-triggered flexibility reduces required positive grid capacity by 2.3–19.6%, while reductions for negative grid load situations range from 9.0–15.6%. These reductions translate into annual incentive budgets of approximately €6,900–31,700 under the forecast-based configuration, corresponding to only €5.3–24.5 per household per year. Under the perfect-information benchmark, in which all critical events are successfully anticipated and mitigated, economically viable incentive budgets increase to €44,300–169,400 annually (€34.2–131.1 per household). Thus, uncertainty reduces economically viable incentive budgets by roughly a factor of three to seven. Even under perfect information, however, economically justifiable incentives remain modest, increasing from below 1 to at most approximately 3,ct per kW shifted, because the larger avoided grid expansion costs are distributed across a substantially greater volume of shifted load.

These incentive levels are modest when compared with evidence from the flexibility literature. At typical Austrian residential electricity prices of around 20–22 ct/kWh in the year of the data

collection 2018, incentives of roughly 1 ct per kW (assuming the shifted capacity is sustained for one hour, corresponding to 1 kWh of shifted electricity) correspond to a price change of around 5%. Combined with commonly reported short-run demand elasticities of approximately −0.1 (Kapeller et al., 2023; Hirth et al., 2022; Labandeira et al., 2017; Faruqui and Sergici, 2010), such signals would imply demand reductions of only around 1–3%. However, these elasticity estimates largely reflect settings requiring repeated manual responses to changing price signals. Emerging home energy management systems may lower this behavioural barrier by shifting the relevant decision from a willingness to repeatedly *act* toward a willingness to *accept* automated control. Evidence from electric vehicle charging supports this interpretation, showing that price differentials of only a few cents per kWh (no directly comparable kW-based estimates exist in the literature) can induce substantial behavioural adjustments once automation is available (Steren et al., 2025; La Nauze et al., 2024).

Several limitations should be acknowledged. The empirical application is illustrative and should not be interpreted as directly transferable to other grid sections. Practical implementation requires location-specific information on network topology, asset constraints and expansion costs. Likewise, threshold definitions, forecasting models and cost assumptions must be tailored to local operating conditions and planning criteria.

More broadly, the framework demonstrates that the value of flexibility is determined by residual peak capacity rather than by shifted energy volumes. Even substantial use of flexibility measures such as load-shifting cannot eliminate the need for grid reinforcement if critical peak events remain unmitigated. Conversely, every avoided residual peak translates directly into avoided infrastructure costs and increases economically justifiable incentive payments. By linking uncertainty, physical system constraints and economic decision-making within a common capacity-based framework, this work provides a transparent basis for evaluating incentive-based flexibility alongside conventional grid expansion as complementary strategies for delivering a cost-effective energy transition.

4. Methods

4.1. Scenario-based load trajectories

PV production in the exemplary village: For residential PV production, we retrieve data from the Renewables Ninja platform (`https://www.renewables.ninja/`) (Renewables Ninja, 2025). By providing geographical location, PV panel size, tilt, orientation and azimuth, detailed hourly production profiles are generated. These are interpolated to 15-minute intervals and cover October 2017 to September 2018. For future adoption scenarios, additional PV production is simulated using standardized PV production profiles. A general panel capacity of 5 kWp is assumed, following the Austrian distribution network study (Schwalbe et al., 2024). The panel orientation distribution is set to 20% south-east-facing, 50% south-facing and 30% south-west-facing, with a fixed tilt angle of 35° and no tracking. Across the scenarios, we assume PV adoption rates of 25.5%, 41.5% and 75%, representing the share of households equipped with PV systems (Schwalbe et al., 2024). To generate the overall PV production at village-level, these household PV production profiles are scaled by the respective adoption rates and multiplied by the number of households in the village to estimate the total PV output.

EV charging profiles in the exemplary village: EV-distance profiles provide information on the distance driven by an electric vehicle between home charging events. These profiles contain infor-

mation on the covered distance for an EV since the last time it was at the respective home-charging-station, assuming no external charging, and are available for January–December 2022 at an hourly resolution. Since the EV profiles do not overlap with the baseline load period (2017–2018), temporal alignment is required. The months October–December are shifted to the start of the year, and minor imputed "gap days" are added around October and the artificial year transition to ensure weekday alignment.

The distance-based EV profiles are then transformed into charging profiles, as the raw data reflect only vehicle movement and not electricity consumption. First, the distance travelled in each interval is converted into energy demand using an average consumption factor of 0.21 kWh per kilometer, consistent with values reported in the literature (Weiss et al., 2024). Since only hourly data is available, these hourly profiles describing the times at which a charging event for the respectively required kWh starts, are first transformed to 15-minute intervals. To avoid charging always starting at the top of the hour, the start time of each charging event is randomly allocated to one of the 15-minute intervals within the respective hour. Second, the resulting energy demand is translated into a charging load by distributing the required kWh across a home-charging event. We assume a maximum charging power of 5.5 kW, following assumptions made by (Schwalbe et al., 2024).

In order to estimate village-level EV consumption, the selected EV charging profiles are scaled by the respective adoption rates (20% or 40%) and multiplied by the number of residents, reflecting that households may own multiple vehicles. This approach assumes each adult can own a vehicle according to the average vehicle-per-resident ratio of 0.643 (Statistik Austria, 2023). Linking EV ownership to the number of residents rather than number of households. When scaling and repeating the available profiles, the consumption of each repeated profile is randomly shifted forward or backward by one or two intervals to avoid unrealistically synchronized charging peaks. Otherwise, peaks would unrealistically increase in magnitude.

Adjustment metric: Having generated village-level PV production and EV charging, the adjusted electricity consumption is computed as:

$$C_{vill_{adj}} = C_{vill_{base}} + EV_{vill} - PV_{vill} \tag{1}$$

4.2. Forecasting

Choice of the Estimation Technique: For implementing incentive-based flexibility, a forecasting model must first be deployed to identify the time intervals in which flexibility interventions is expected to be necessary. The model should provide forecasts of the expected load level for each 15-minute interval of the following day while also the probabilities that the adjusted load exceeds predefined upper thresholds or falls below zero are of interest. A modelling framework only returning binary event classifications would be insufficient, as it identifies whether a threshold event is expected to occur but does not directly provide information on the magnitude of the underlying load which would need to be shifted.

Several modelling approaches are available for such forecasting tasks. Regression models offer high interpretability but may be too restrictive when the relationship between load and weather conditions is non-linear (Wang, 2020). More flexible non-parametric approaches, such as Random

Forests, Gradient Boosting methods, Support Vector Machines and neural-network-based models, can capture non-linearities and interaction effects more effectively (O'Brien and Ishwaran, 2019; Rosales-Pérez et al., 2023; Scheepens et al., 2023). Among the available forecasting methods, gradient boosted quantile regression provides a suitable framework for the task at hand as it estimates conditional quantiles of the response variable rather than only the conditional mean (Koenker, 2005). This is particularly useful in the present setting, where the objective is not only to forecast the continuous load time series, but also to derive probabilities for predefined threshold events. Estimating these quantities with separate models could lead to inconsistencies between the predicted load trajectory and the associated event probabilities. For instance, a classification model may assign a high probability to a threshold exceedance, while an independently estimated point forecast remains below the same threshold. The quantile-based approach avoids this issue by deriving both the central load forecast and the exceedance probabilities from a single predictive distribution.

Selected Model: In this paper, the conditional quantiles are estimated using gradient boosting regression with quantile loss. Gradient boosting is well suited to the present application because it can capture nonlinear relationships and interactions between lagged load, weather variables and temporal features, while remaining less computationally demanding than many deep-learning alternatives. For each 15-minute interval of the day-ahead horizon, the model estimates quantile forecasts for the levels

$$\alpha \in \{0.01, 0.05, 0.10, 0.25, 0.50, 0.75, 0.90, 0.95, 0.99\}$$

The median quantile is used as the central point forecast of adjusted load. The full set of estimated quantiles provides a discrete approximation of the conditional predictive distribution and is used to derive the threshold-event probabilities considered in the analysis.

For a given threshold q, the exceedance probability is approximated as

$$P(C_{vill_{adj},i} > q) \tag{2}$$

where $C_{vill_{adj}}$ denotes the adjusted load of the district in interval i, with $i \in \{1, 2, \ldots, 96\}$ representing the 96 15-minute intervals of the day-ahead forecast horizon. The probability is obtained by locating q_{lim} within the predicted quantile function and interpolating between neighbouring quantile estimates. Analogously, the probability of negative load situations is obtained by evaluating the predicted quantile function at zero. To ensure a valid ordering of the estimated distribution, the predicted quantiles are post-processed such that lower quantiles cannot exceed higher quantiles.

Training process: The training process for the forecasting model follows a sequential walk-forward approach. Predictor variables comprise lagged observations of the adjusted load, day-ahead forecasted climate variables, and engineered temporal features, as detailed later in this Section. Based on these inputs, the model generates day-ahead forecasts for each 15-minute interval of the subsequent 24-hour horizon. Model performance is evaluated over an out-of-sample period spanning October 2017 to September 2018. Forecasting proceeds in a temporally ordered manner: for each target month, the model is trained using observations from the preceding month. This design ensures that each forecast is conditioned only on information available prior to the forecast period and therefore preserves temporal causality. For each target month, gradient boosting models with

quantile loss are estimated separately for the different quantile levels. Hyperparameters are optimized using a time series cross-validation procedure based on rolling-origin validation, where each split trains the model on past observations and evaluates it on the immediately subsequent validation window, preventing leakage of future information into the tuning process. Hyperparameters are selected based on the pinball loss of the median quantile model and are subsequently applied to the remaining quantile models. The resulting set of quantile forecasts is then used to obtain the median load forecast, prediction intervals, and the threshold-event probabilities used in the subsequent analysis.

Feature selection: To support the research question and subsequent modelling, meteorological data is incorporated as weather strongly influences consumption pattern and load demands. Historic forecasted temperature and global radiation data for each following day is obtained from BlueSky Wetteranalysen (`https://www.blueskywetter.com`) (BlueSky Wetteranalysen, 2018), initially collected for the Austrian electricity provider LinzAG (`https://www.linzag.at`) (LinzAG, 2018) and is available on an hourly resolution for the period of July 2017 to September 2018. To ensure consistency with the 15-minute resolution of the remaining data, the hourly forecasts were interpolated to 15-minute intervals. These historical forecast data represent the as-issued day-ahead forecasts available at the time for each corresponding day in 2017 and 2018. They therefore constitute a faithful record of the information that would have been accessible to the model in real time, rather than actual observed or re-analysis data. In addition to these meteorological inputs, a range of time-related calendar features are engineered to improve modelling performance. These include but are not limited to sinusoidal transformations to approximate the time-of-day pattern, indicator variables for weekends and public holidays, weekday-specific effects as well as variables capturing whether a given interval fell within standard working hours alongside lagged observations of the load itself

4.3. Annualized investment budget

Grid investments are assumed to be recovered over a depreciation horizon of 20 years, consistent with the regulatory depreciation periods for low-voltage transformers and cables specified in the Austrian income tax guidelines issued by the Federal Ministry of Finance (Bundesministerium für Finanzen, 2024). To convert upfront investment costs into an equivalent annual budget, we apply a standard annuity formula using a discount rate of 6.11%, corresponding to the weighted average cost of capital (WACC) set by the Austrian energy regulator E-Control for long-term electricity distribution investments (E-Control Austria, 2025). The use of the regulated WACC is appropriate in this context, as it reflects the allowed rate of return for regulated grid operators and therefore represents the relevant opportunity cost of capital for grid expansion decisions.

The resulting annualized budget represents the maximum amount available for incentive-based flexibility. To derive an indicative incentive level per unit of shifted load, this annual budget is divided by the total volume of forecast-identified threshold events across all intervals. These intervals correspond to cases in which the forecasting model signals that an upper threshold exceedance, or a negative-load situation, is sufficiently likely and therefore triggers a flexibility intervention, including all triggered interventions, irrespective of whether the threshold event subsequently materializes. This reflects the fact that financial incentives would already have been offered whenever the model issued a signal, including in cases that later turn out to be false alarms. Dividing the annualized budget by the total incentivized load volume then yields an implied incentive payment per unit of shifted load, which can be interpreted as an order-of-magnitude benchmark for eco-

nomically viable incentive payments under the assumed conditions. In the empirical application, this benchmark is reported in capacity-based terms, consistent with the fact that grid expansion costs are driven by peak capacity requirements rather than annual electricity volumes.

This exercise relies on a set of strong simplifying assumptions intended to illustrate an upper bound on viable incentive levels. We take the following assumptions: (i) threshold exceedances are assumed to be known sufficiently in advance, incentives are assumed to be immediately accepted by households, and electricity consumption is assumed to be shifted without delay; (ii) all threshold exceedances are treated equally, i.e., the probability of an exceedance is not explicitly weighted in the budgeting exercise; (iii) the exceedance electricity is shifted fully to non-exceedance times.

Data availability

The replication package contains all data to reproduce the presented analyses and results. It includes the cleaned and imputed household electricity consumption, electric vehicle demand and photovoltaic electricity generation under the assumed adoption rates at an aggregated level, as well as the explanatory variables used in the analysis. The anonymized raw load profile data of the 1292 households at household-level in our study is available from the authors upon reasonable request.

Acknowledgements

We gratefully acknowledge funding from the Austrian Federal Ministry for Climate Action, Environment, Energy, Mobility, Innovation and Technology for Valentin Praun and Katharina Rusch through the Powercast project (grant agreement no. F0999910227), and from the Austrian Federal Ministry of Innovation, Mobility and Infrastructure for Johannes Reichl through the Activate project (grant no. F099943528). The funders had no role in the study design, data collection, analysis and interpretation of the data, or the writing of the manuscript.

References

Agency for the Cooperation of Energy Regulators (ACER), 2023. Market monitoring report 2023: Barriers to demand response.

Albadi, M.H., El-Saadany, E.F., 2007. Demand response in electricity markets: An overview, in: 2007 IEEE power engineering society general meeting, IEEE. pp. 1–5. doi:https://doi.org/10.1109/PES.2007.385728.

Austrian Power Grid (APG), 2025. 9 Milliarden Euro für eine versorgungssichere Energiewende. Accessed May 2025.

Azarova, V., Engel, D., Ferner, C., Kollmann, A., Reichl, J., 2018. Exploring the impact of network tariffs on household electricity expenditures using load profiles and socio-economic characteristics. Nature Energy 3, 317–325. doi:10.1038/s41560-018-0105-4.

BlueSky Wetteranalysen, 2018. BlueSky Wetteranalysen. https://www.blueskywetter.com/. Accessed June 17, 2026.

Borenstein, S., 2005. The long-run efficiency of real-time electricity pricing. The Energy Journal 26, 93–116. doi:10.5547/ISSN0195-6574-EJ-Vol26-No3-5.

Bradley, P., Coke, A., Leach, M., 2016. Financial incentive approaches for reducing peak electricity demand, experience from pilot trials with a uk energy provider. Energy Policy 98, 108–120.

Brown, T., Schlachtberger, D., Kies, A., Schramm, S., Greiner, M., 2018. Synergies of sector coupling and transmission reinforcement in a cost-optimised, highly renewable european energy system. Energy 160, 720–739. doi:10.1016/j.energy.2018.08.121.

Bundesministerium für Finanzen, 2024. Einkommensteuerrichtlinien 2000 (EStR 2000) – Abschreibung von Anlagevermögen.

Burger, S.P., Luke, M., 2017. Business models for distributed energy resources: A review and empirical analysis. Energy Policy 109, 230–248. doi:10.1016/j.enpol.2017.07.007.

Council of European Energy Regulators, 2019. Italy nr – power level choices for domestic consumers. https://www.ceer.eu/wp-content/uploads/2024/04/C19_NR_Italy_NL.pdf. Describes the ability for households to select finer-grained contracted power levels and the operational consequences of exceeding them.

E-Control Austria, 2024. Technical and Organizational Rules for Operators and Users of Networks (TOR). Accessed: 2025-04-07.

E-Control Austria, 2025. WACC für Stromverteilung, Stromübertragung und Gasverteilung – Neuinvestitionen.

European Commission, 2019. Clean energy for all europeans.

European Commission, 2020. An eu strategy for energy system integration.

European Commission, 2024. Convergence report 2024.

Faruqui, A., Sergici, S., 2010. Household response to dynamic pricing of electricity: A survey of evidence. Energy Journal 31, 193–225.

Hernandez Arias, V.A., Arias-Barragán, L.A., Rivas-Trujillo, E., 2020. Incentive-based demand response: Case study. Scientia et Technica 25, 216–222.

Hirth, L., Khanna, T., Ruhnau, O., 2022. The (very) short-term price elasticity of German electricity demand. Technical Report. URL: https://hdl.handle.net/10419/249570.

Hoffmann, P., Krichene, H., Zimmer, M., Mier, M., 2025. Plug, baby, plug: Unlocking europe's electricity market. Special Report of Allianz Research, Munich. Accessed: 2025-04-07.

Hofmann, M., Bjarghov, S., Sæle, H., Lindberg, K.B., 2025. Grid tariff design and peak demand shaving: A comparative tariff analysis with simulated demand response. Energy Policy 198, 114475. doi:10.1016/j.enpol.2024.114475.

Intergovernmental Panel on Climate Change, 2022. Climate Change 2022: Mitigation of Climate Change. Cambridge University Press. doi:10.1017/9781009157926.

International Energy Agency, 2023. Electricity Grids and Secure Energy Transitions. doi:10.1787/6c7f2c9a-en.

Ito, K., Ida, T., Tanaka, M., 2023. Selection on welfare gains: Experimental evidence from electricity plan choice. American Economic Review 113, 2937–2973. doi:10.1257/aer.20210150.

Joskow, P.L., 2007. Competitive Electricity Markets and Investment in New Generating Capacity, in: Sioshansi, F.P. (Ed.), The New Energy Paradigm. Elsevier, pp. 76–122. doi:10.1016/B978-012375628-2/50006-2.

Kapeller, R., Cohen, J.J., Kollmann, A., Reichl, J., 2023. Incentivizing residential electricity consumers to increase demand during periods of high local solar generation. Energy Economics 127, 107028.

Koenker, R., 2005. Quantile Regression. Econometric Society Monographs, Cambridge University Press, Cambridge.

La Nauze, A., Friesen, L., Lim, K.L., Menezes, F., Page, L., Philip, T., Whitehead, J., 2024. Can electric vehicles aid the renewable transition? evidence from a field experiment incentivising midday charging. Technical Report. CESifo Working Paper.

Labandeira, X., Labeaga, J.M., López-Otero, X., 2017. A meta-analysis on the price elasticity of energy demand. Energy Policy 102, 549–568. doi:10.1016/j.enpol.2017.01.002.

LinzAG, 2018. LinzAG. https://www.linzag.at/. Accessed June 17, 2026.

Lopes, J.A.P., Hatziargyriou, N., Mutale, J., Djapic, P., Jenkins, N., 2007. Integrating distributed generation into electric power systems: A review of drivers, challenges and opportunities. Electric Power Systems Research 77, 1189–1203. doi:10.1016/j.epsr.2006.08.016.

O'Brien, R., Ishwaran, H., 2019. A random forests quantile classifier for class imbalanced data. Pattern Recognition 90, 232–249. doi:10.1016/j.patcog.2019.01.036.

Oesterreichs Energie, 2025. Die Entwicklung der Netzkosten und Netztarife in Österreich.

Palensky, P., Dietrich, D., 2011. Demand side management: Demand response, intelligent energy systems, and smart loads. IEEE Transactions on Industrial Informatics 7, 381–388. doi:10.1109/TII.2011.2158841.

Renewables Ninja, 2025. Renewables.ninja: Hourly solar & wind power simulations. Accessed June 17, 2025.

Rosales-Pérez, A., García, S., Herrera, F., 2023. Handling Imbalanced Classification Problems With Support Vector Machines via Evolutionary Bilevel Optimization. IEEE Transactions on Cybernetics 53, 4735–4747. doi:10.1109/tcyb.2022.3163974.

Ruhnau, O., Hennig, P., Madlener, R., 2020. Economic implications of forecasting electricity generation from variable renewable energy sources. Renewable Energy 161, 1318–1327. doi:10.1016/j.renene.2020.06.110.

Scheepens, D.R., Schicker, I., Hlaváčková-Schindler, K., Plant, C., 2023. Adapting a deep convolutional rnn model with imbalanced regression loss for improved spatio-temporal forecasting of extreme wind speed events in the short to medium range. Geoscientific Model Development 16, 251–270. doi:10.5194/gmd-16-251-2023.

Schwalbe, R., Brunner, H., Scheida, K., Brandl, S., Groiß, C., Leitner, M., Schoppitsch, D., Sporer, W., Wieland, T., 2024. Aktualisierung der netzberechnungen der studie „volkswirtschaftlicher wert der stromverteilernetze auf dem weg zur klimaneutralität in Österreich“.

Statistik Austria, 2023. Statistik der Kraftfahrzeuge 2022. Accessed: 2025-02-15.

Steiner, P.O., 1957. Peak loads and efficient pricing. Quarterly Journal of Economics 71, 585–610. doi:10.2307/1885679.

Steren, A., Parag, Y., Teschner, N., Zemah-Shamir, S., 2025. Consumer preferences and willingness to pay for ev charging: Implications for incentives to promote off-peak charging and renewables integration. Applied Energy 392, 125981. doi:10.1016/j.apenergy.2025.125981.

Turitsyn, K., Sulc, P., Backhaus, S., Chertkov, M., 2011. Options for control of reactive power by distributed photovoltaic generators. Proceedings of the IEEE 99, 1063–1073. doi:10.1109/JPROC.2011.2116750.

U.S. Department of Energy, 2006. Benefits of Demand Response in Electricity Markets and Recommendations for Achieving Them: Report to Congress pursuant to Section 1252 of the Energy Policy Act of 2005. Technical Report. U.S. Department of Energy, Office of Electricity Delivery and Energy Reliability. URL: https://www.energy.gov/sites/prod/files/oeprod/DocumentsandMedia/DOE_Benefits_of_Demand_Response_in_Electricity_Markets_and_Recommendations_for_Achieving_Them_Report_to_Congress.pdf.

Wang, H., 2020. Logistic Regression for Massive Data with Rare Events, in: III, H.D., Singh, A. (Eds.), Proceedings of the 37th International Conference on Machine Learning, PMLR. pp. 9829–9836.

Weiss, M., Newman, A., Helmers, E., 2024. Energy Consumption of Electric Vehicles in Europe. Sustainability 16. doi:10.3390/su16177529.

Weitzman, M.L., 1974. Prices vs. quantities. The Review of Economic Studies 41, 477–491. doi:10.2307/2296698.

Wolak, F.A., 2007. Residential Customer Response to Real-time Pricing: The Anaheim Critical Peak Pricing Experiment. Technical Report. Center for the Study of Energy Markets, UC Berkeley. URL: https://escholarship.org/uc/item/3td3n1x1.

Appendix

A Motivation for incentive-based flexibility and household requirements 23
B Welfare economic rationale 24
C Imputation of baseline electricity consumption 26
D Descriptive statistics of adjusted load 27
E Modelling results: distribution under Q_{85} threshold 28
F Scenario-specific post-intervention capacity requirements 28
G Economic valuation of forecast-based flexibility 29

This appendix provides additional tables, figures, and material discussed in the manuscript "The economic value of load flexibility and its viability as a substitute for grid expansion" by Valentin Praun, Johannes Reichl, and Katharina Rusch.

A. Motivation for incentive-based flexibility and household requirements

To mitigate critical load situations and reduce the need for costly grid upgrades, the literature discusses two broad approaches to smoothing the LDC: (i) technological or quantity-based measures, which physically restrict consumption or feed-in, and (ii) economic or incentive-based measures, which influence consumer behavior through price signals.

Technological measures impose limits on electricity consumption directly. For instance, in Italy, low-voltage electricity customers have a contracted maximum connection capacity (*potenza impegnata*), typically 3 kW. Smart meters allow short-term peaks of around 10% above this value (*potenza disponibile*, approximately 3.3 kW); if this limit is exceeded, the meter's internal limiter disconnects the supply, which can be restored manually after reducing the load (see e.g. Council of European Energy Regulators, 2019). While such mechanisms can restrict individual peak demand, the load of a single household is generally small relative to the aggregated demand within a grid section. Consequently, disconnecting individual consumers is typically neither necessary nor effective for smoothing peaks in the LDC. More targeted technological solutions would require monitoring aggregated load within a grid area and initiating curtailment only when a predefined load threshold is close to being reached.

Economic approaches instead rely on incentives that encourage households to adjust electricity consumption during periods of critical system stress (Hernandez Arias et al., 2020; Hirth et al., 2022; Burger and Luke, 2017; Agency for the Cooperation of Energy Regulators (ACER), 2023; European Commission, 2024; Labandeira et al., 2017; Hofmann et al., 2025). While grid congestion and reinforcement costs are ultimately recovered from consumers through network tariffs, these charges are typically based on ex post average-cost recovery. Households therefore do not face the marginal, local, and time-specific cost of consumption during periods of acute grid stress. Targeted demand-response incentives can help internalize these congestion-related externalities by rewarding households for reducing or shifting demand when grid capacity is scarce.

Compared with fixed technical limits, incentive-based schemes preserve consumer choice and can improve allocative efficiency. Household electricity uses differ in both (Willingness-to-pay (WTP)) and flexibility. Loads such as EV charging, dishwashers, or laundry are often more elastic and easier to shift, whereas uses such as lighting or cooking at fixed times are less responsive. By

offering monetary rewards for temporary load reductions, demand-response schemes can target flexible, lower-value loads while leaving higher-value, less-flexible consumption largely unaffected. Existing experimental studies document household responses to price signals and incentive-based flexibility (Kapeller et al., 2023; Bradley et al., 2016; Wolak, 2007; Ito et al., 2023), although incentive levels are often set exogenously rather than derived from avoided grid investment costs.

From a welfare-economic perspective, such schemes can be interpreted as second-best instruments. They do not achieve the allocative precision of fully efficient real-time marginal-cost pricing, but they may approximate it when first-best pricing is infeasible (Weitzman, 1974). Building on the heterogeneity in WTP and demand elasticity, Weitzman's framework implies that quantity-based controls are more effective when demand is steep, whereas price-based instruments perform better when demand is more elastic. Empirical evidence points to substantial heterogeneity in household demand elasticities (U.S. Department of Energy, 2006), and Ito et al. (2023) show that more elastic consumers are more likely to self-select into demand-response programs, supporting the potential welfare advantages of incentive-based approaches.

The practical success of incentive-based flexibility depends on accurate and timely signals. Unlike automated technical curtailment, behavioural responses require advance information, since households need time to plan flexible consumption. Heat pumps, for example, can provide flexibility but should not be disconnected abruptly, as this may disrupt heating cycles or damage equipment. Incentives must therefore often be communicated several hours before a critical load period, which requires reliable forecasts of aggregated load at the transformer or grid-section level. At the same time, ex-ante communication of congestion forecasts or incentives may create feedback effects: households may adapt their behaviour in anticipation of expected incentives, thereby affecting the realized load and potentially the validity of the original forecast. This highlights the importance of carefully designing the forecasting and communication mechanism.

An increasing share of households rely on home automation systems, including smart heat pumps, battery management systems, and automated EV chargers, which can autonomously shift flexible loads or allocate surplus PV generation without active consumer intervention. Such automation can enhance responsiveness and help ensure that flexibility is deployed reliably and repeatedly when incentives are triggered.

B. Welfare economic rationale

The framework developed in this paper derives an upper bound for economically viable flexibility incentives from avoided grid expansion costs. This appendix provides the welfare-economic rationale behind this approach. The central argument is that local peak loads create system-level costs that are not reflected in the prices faced by individual consumers. As a result, decentralized consumption decisions can be privately optimal while still leading to socially inefficient peak loads. Incentive-based flexibility can correct this missing price signal, but only up to the point at which the cost of incentives does not exceed the avoided cost of grid expansion.

In an idealized competitive market with complete price signals, consumers would face prices that reflect the marginal system cost of electricity consumption at each point in time and location. Under standard assumptions, such a competitive equilibrium maximizes *total surplus*, defined as the overall economic value generated in the system: the sum of consumer benefits from electricity consumption minus the cost of providing it. In real-world distribution grids, however, households typically face electricity tariffs that are uniform across broad regions and largely independent of

local transformer load. Consumers therefore choose electricity consumption by comparing their marginal willingness to pay with the electricity tariff they face, but this tariff does not reflect the additional grid expansion costs caused by consumption during local peak periods.

This creates a wedge between private and social costs. From the individual household perspective, consuming electricity during a peak period may be optimal because the price does not increase with local grid stress. From the system perspective, however, the same consumption may contribute to higher peak capacity requirements and therefore to costly grid reinforcement. Peak-induced grid expansion costs are thus not internalized by individual consumers. The resulting allocation can therefore fail to maximize social welfare, even if consumers behave rationally given the prices they face. The role of flexibility incentives is to approximate the missing scarcity signal and induce demand reductions during periods in which the system-level value of avoided peak capacity is high.

A simple alternative to incentive-based flexibility would be a uniform consumption cap during critical periods. Such caps can reduce peak demand, but they generally do so inefficiently because they ignore heterogeneity in consumers' willingness to pay. A uniform cap restricts all consumers equally, regardless of whether the restricted electricity use has high or low value for the respective household. This can generate *deadweight losses*, i.e. economic value that is lost because electricity is not consumed by those who value it most. For example, if one consumer has a high willingness to pay for electricity use during a critical period but is restricted by a cap, while another consumer with lower willingness to pay consumes within the same limit, potential surplus is lost. In contrast, a monetary incentive allows consumers to self-select into reducing demand where the private value of consumption is relatively low.

Formally, consider N consumers indexed by $i = 1, \dots, N$, each with quasilinear preferences and inverse demand, or marginal willingness to pay, $v_i(q_i)$ for electricity consumption q_i. Each v_i is continuous and strictly decreasing, reflecting that additional units of electricity are valued less than earlier ones. Let aggregate consumption be

$$Q = \sum_{i=1}^{N} q_i,$$

and denote the marginal cost of electricity supply as $c(Q)$, which is continuous and weakly increasing.

The *total surplus* TS of an allocation $(q_1, \dots, q_N)$ is the difference between the total value that consumers derive from consuming electricity and the total cost of providing it:

$$TS = \underbrace{\sum_{i=1}^{N} \int_0^{q_i} v_i(x)\, dx}_{\text{total consumer value}} - \underbrace{\int_0^{Q} c(x)\, dx}_{\text{total supply cost}}.$$

In a benchmark allocation with complete marginal cost pricing, consumption satisfies

$$v_i(q_i^*) = c(Q^*) \quad \text{for all } i \text{ with } q_i^* > 0, \qquad \sum_{i=1}^{N} q_i^* = Q^*.$$

This condition implies that marginal willingness to pay is equalized across consumers and aligned with marginal system cost. No unit of electricity is allocated to a consumer with lower marginal value while another consumer with higher marginal value is rationed.

Now suppose a policymaker introduces a uniform consumption cap $\bar{q}$, restricting each consumer to

$$q_i \leq \bar{q} \quad \text{for all } i.$$

If all consumers are bound by the cap, aggregate consumption becomes

$$Q_{\bar{q}} = N\bar{q},$$

and each consumer consumes exactly $\bar{q}$. Because consumers are heterogeneous, their marginal valuations at $\bar{q}$ generally differ:

$$v_i(\bar{q}) \neq v_j(\bar{q}).$$

Some consumers are prevented from consuming units they value highly, while others continue to consume units they value less. The cap therefore violates the equimarginal principle of welfare economics and cannot maximize total surplus, even conditional on the same aggregate quantity $Q_{\bar{q}}$. In the context of flexibility, this means that a uniform restriction may reduce peak load, but it does not ensure that the reduction occurs where it is least costly for consumers.

An incentive scheme provides a more efficient alternative because it allows consumers to reveal, through their response, whether reducing consumption is privately costly or relatively easy. Suppose each consumer receives a monetary reward $t > 0$ per unit of reduced load or shifted consumption during a critical period, measured relative to a predefined baseline. Economically, this is equivalent to increasing the opportunity cost of consuming electricity during that period. Let $(q_1^t, \ldots, q_N^t)$ denote the allocation under the incentive scheme. The allocation satisfies

$$v_i(q_i^t) = c(Q_t) + t \quad \text{for all } i \text{ with } q_i^t > 0, \qquad \sum_{i=1}^{N} q_i^t = Q_t,$$

where $Q_t = \sum_i q_i^t$ is aggregate consumption under the incentive.

The incentive increases the effective marginal cost of consuming electricity during the critical period. Consumers whose marginal value of consumption is below this effective cost reduce their demand, while consumers with high marginal value continue to consume. Hence, total consumption during the critical period falls relative to the benchmark without imposing identical quantity reductions on all users. Importantly, the incentive scheme preserves the allocation of scarce grid capacity to those consumers who value electricity most. It therefore reduces peak demand in a way that is less welfare-reducing than a uniform cap.

C. Imputation of baseline electricity consumption

The baseline load profile was constructed from household-level electricity consumption data recorded at 15-minute intervals. Since consumption data were not available for all households in every interval, the observed aggregate load was adjusted for missing observations.

A simple correction based only on the share of missing households was not considered appropriate, because households differ substantially in size and electricity consumption. In particular, the omission of a large household is likely to have a greater effect on the aggregate load than the omission of a small household. Residential floor area was therefore used to account for these differences.

Let $C_{h,t}$ denote the observed electricity consumption of household h in interval t, and let A_h denote its residential floor area. For each interval, the observed aggregate baseline consumption was calculated as

$$C^{\text{obs}}_{\text{vill,base},t} = \sum_{h \in O_t} C_{h,t}, \tag{C.1}$$

where O_t is the set of households with an available observation in interval t. The share of the total residential floor area represented by these households was calculated as

$$s_t = \frac{\sum_{h \in O_t} A_h}{\sum_h A_h}. \tag{C.2}$$

The aggregate baseline consumption was then adjusted according to

$$\widehat{C}_{\text{vill,base},t} = \frac{C^{\text{obs}}_{\text{vill,base},t}}{s_t}. \tag{C.3}$$

The observed aggregate baseline load is therefore scaled by the share of total residential floor area covered in each interval. It assumes that the average electricity consumption per unit of floor area is similar for households with available and missing observations.

The resulting series was reindexed to obtain a complete sequence of 15-minute intervals. Five days were entirely missing from the time series. These days were replaced by the load profile of a manually selected comparable day with similar calendar and seasonal characteristics. In addition, three implausible individual observations were identified and replaced by the value recorded 30 minutes earlier.

D. Descriptive statistics of adjusted load

For completeness, Table D.3 summarizes the distribution of adjusted electrical load across the three simulated EV and PV penetration scenarios. Negative values indicate periods in which local PV production exceeds electricity demand, resulting in net power export, whereas positive values represent net electricity consumption.

Table D.3 Summary statistics of adjusted load for all simulated EV and PV adoption scenarios. (Values in kW)

Statistic	**EV 20%, PV 25.5%**	**EV 40%, PV 41.5%**	**EV 75%, PV 75%**
Minimum	-705.1	-1479.0	-3101.3
10th percentile	-167.8	-632.74	-1628.7
25th percentile	266.1	129.9	-221.0
Median	445.4	483.1	541.3
Mean	452.8	421.7	316.1
75th percentile	733.5	851.8	1071.5
90th percentile	1005.0	1243.0	1664.2
Maximum	1580.5	2045.8	2885.3

E. Modelling results: distribution under Q_{85} threshold

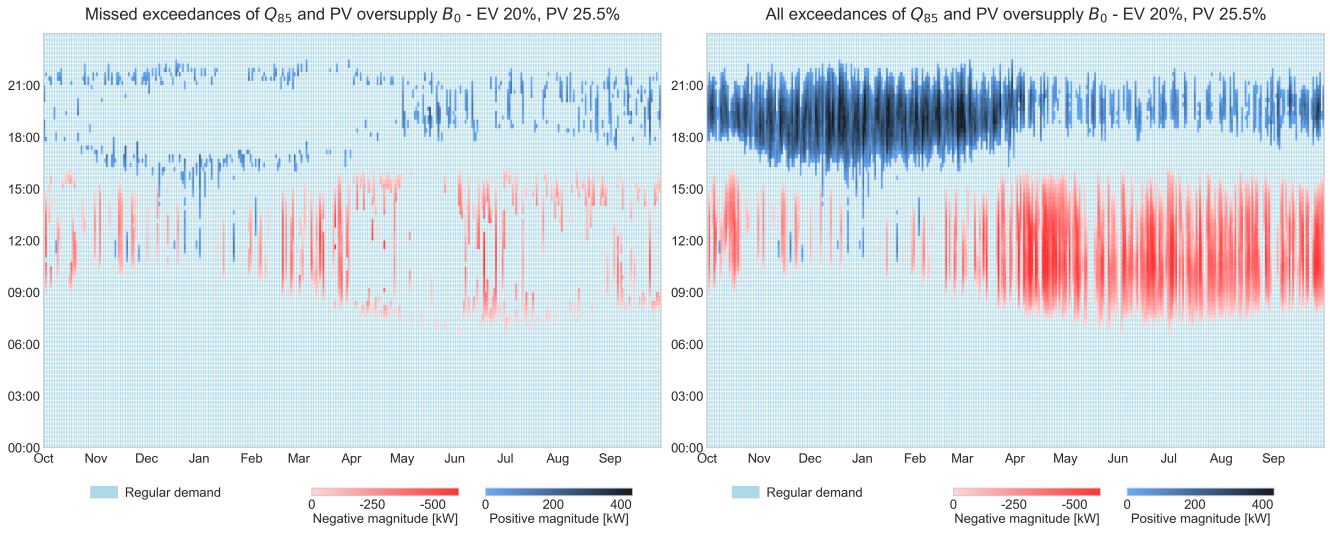


Figure E.7 - Classification results false negatives only (EV 20%, PV 25.5%) under the Q_{85} threshold: The left panel shows all time intervals that the model incorrectly classified as non-events (false negatives): dark blue marks missed exceedances above Q_{85}, while red marks missed periods of PV oversupply. The right panel shows the corresponding true exceedances for comparison. To maintain the focus on the more critical false negatives for now, any false positives are shaded in light blue as if they were "normal" non-events. The false positive frequency is shown in Figure 6. Overall, the model captures the majority of actual exceedances and oversupply periods. Most missed events which are left occur near the boundaries of exceedance hours for example, when the model predicts the event to begin slightly too late or end slightly too early, apart from a few days at the beginning of the forecasting window where sometimes, entire days were missed. which occurs primarily at the beginning of the forecasting window, where the model undergoes its strongest parameter adaptation and is constrained by a sparse training dataset.

F. Scenario-specific post-intervention capacity requirements

While in the main text only the key findings regarding the post-intervention capacity requirements were summarised, this appendix provides a more detailed description of the capacity reductions shown in Figure 6. The Figure compares the pre- and post-intervention capacity requirements after targeting peak-load events for the three considered EV and PV adoption scenarios. In all scenarios, the upper threshold for positive peak-load events is varied between the Q_{85} and Q_{95} quantiles, whereas the threshold for negative-load events is kept constant at zero. Across all scenarios, the model-based intervention reduces the required capacity range relative to the pre-intervention load trajectories.

For the lowest adoption scenario, (20% EV, 25.5% PV), the pre-intervention capacity requirement ranges from -705 kW to 1580 kW. Under the Q_{85} threshold, where positive peaks above 950 kW and negative-load events are targeted, the post-intervention range is reduced to -642 kW to 1270 kW. This corresponds to an overall capacity-range reduction of 16.3%, with reductions of 19.6% at the upper bound and 9.0% at the lower bound. Under the Q_{95} threshold, targeting positive peaks above 1147 kW and negative-load events, the post-intervention capacity requirement ranges from -642 kW to 1467 kW, yielding an overall reduction of 5.0%, consisting of reductions of 7.2% at the upper bound and 9.0% at the lower bound.

For the intermediate adoption scenario, (40% EV, 41.5% PV), the pre-intervention capacity requirement ranges from −1478 kW to 2045 kW. Under the Q_{85} threshold, targeting positive peaks above 1195 kW and negative-load events, the post-intervention requirement is reduced to −1247 kW to 1716 kW. This corresponds to an overall reduction of 15.9%, with reductions of 16.1% at the upper bound and 15.6% at the lower bound. Under the Q_{95} threshold, where positive peaks above 1442 kW and negative-load events are targeted, the post-intervention range is −1247 kW to 1863 kW, representing an overall reduction of 11.7%, with reductions of 8.9% at the upper bound and 15.6% at the lower bound.

For the highest adoption scenario, (75% EV, 75% PV), the pre-intervention capacity requirement increases substantially, spanning from −3101 kW to 2885 kW. Under the Q_{85} threshold, targeting peaks above 1630 kW and negative-load events, the post-intervention range is reduced to −2713 kW to 2459 kW. This corresponds to an overall capacity-range reduction of 13.6%, comprising reductions of 14.8% at the upper bound and 12.5% at the lower bound. Under the Q_{95} threshold, targeting positive peaks above 2004 kW and negative-load events, the post-intervention capacity requirement ranges from −2713 kW to 2820 kW. This yields an overall reduction of 7.6%, with reductions of 2.3% at the upper bound and 12.5% at the lower bound.

Across all targets, the false positive frequency stays at minimal levels, remaining between 2.0% and 3.7% for negative load conditions, between 1.8% and 2.4% for exceedances above the Q_{85} threshold and between 0.6% and 1.1% for exceedances above the Q_{95} threshold.

G. Economic valuation of forecast-based flexibility

Let L_t denote the observed aggregate load at time step $t \in \{1, \ldots, T\}$, measured at the relevant grid node. For each scenario, we define a critical capacity threshold q, above which load is considered capacity-relevant. This threshold may be derived from the empirical baseline load distribution, for example as a high quantile, or from a scenario-specific technical capacity limit.

The forecasting model provides an estimate of future load, denoted by $\hat{L}_t$. Based on this forecast, we define the requested flexibility volume as the predicted exceedance above the critical threshold:

$$\hat{e}_t = \max(0, \hat{L}_t - q).$$

This quantity represents the amount of flexibility that would be requested by the grid operator at time t. Importantly, $\hat{e}_t$ is a decision variable derived from the forecast, not an observed quantity.

The actual threshold exceedance is given by

$$e_t = \max(0, L_t - q).$$

Since load can only be shifted if an exceedance actually occurs, the effectively shifted volume is bounded by the smaller of the requested and the actual exceedance:

$$e_t^{\text{eff}} = \min(\hat{e}_t, e_t).$$

The resulting post-intervention load profile is then constructed as the counterfactual residual load

$$L_t^{\text{res}} = L_t - e_t^{\text{eff}}.$$

This profile is not observed in the data, but represents an upper-bound counterfactual in which all requested flexibility is fully and instantaneously realised whenever the corresponding load actually exists.

The capacity requirement without flexibility is defined as

$$K^{\text{baseline}} = \max_t L_t,$$

whereas the remaining capacity requirement after forecast-based flexibility is

$$K^{\text{new}} = \max_t L_t^{\text{res}}.$$

The potentially avoidable capacity requirement is therefore

$$\Delta K = K^{\text{baseline}} - K^{\text{new}}.$$

Let C_{grid} denote the unit cost of grid expansion, expressed in € per kW of additional capacity. The avoided grid expansion cost is then

$$C_{\text{grid}} \cdot \Delta K.$$

The total volume of requested flexibility is given by

$$\sum_{t=1}^{T} \hat{e}_t.$$

This quantity includes both correct intervention signals and false-positive intervention signals. Hence, overly aggressive forecasting rules that request flexibility too frequently increase the denominator even if they do not contribute to avoided capacity.

We define the maximum economically justifiable incentive level as

$$p^{\max} = \frac{C_{\text{grid}} \cdot \Delta K}{\sum_{t=1}^{T} \hat{e}_t}.$$

The value $p^{\max}$ can be interpreted as the highest payment per requested unit of shifted electricity for which forecast-based flexibility remains no more expensive than conventional grid expansion.

This formulation penalises both types of forecasting errors. False-negative predictions reduce the avoided capacity ΔK, since critical peaks remain unmitigated and therefore increase K^{new}. False-positive predictions increase the total requested shifting volume $\sum_t \hat{e}_t$, without necessarily reducing the capacity requirement. The proposed metric therefore links forecast quality directly to its economic value for grid planning decisions.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- npjCleanEPaperLoadFlexibility1Model.ipynb
- npjCleanEPaperLoadFlexibility3Figure3.r
- npjCleanEPaperLoadFlexibility2Plots.ipynb
- dfEV40PV415.7z
- dfEV75PV75.7z
- dfEV20PV255.csv
- forecastEV20PV255dayahead.7z
- forecastEV40PV415dayahead.7z
- forecastEV75PV75dayahead.7z