Home Battery Dispatch under a Tiered Peak Power Tariff

David Pérez-Piñeiro* Sigurd Skogestad* Stephen Boyd†

May 5, 2026

Abstract

We consider the problem of operating a battery in a home connected to the grid to minimize electricity cost, which combines an energy charge and a tiered peak power charge based on the average of the N largest daily peak powers in each billing month. With perfect foresight of loads and prices, the minimum cost is the solution of a mixed-integer linear program (MILP), which provides a lower bound on the cost of any implementable policy. We propose a model predictive control (MPC) policy that uses simple forecasts of loads and prices and solves a small MILP at each time step. Numerical experiments on one year of data from a home in Trondheim, Norway, show that the MPC policy attains a cost within 1.7% of the prescient bound, and saves close to three times as much as the best rule-based policy we consider.

*Department of Chemical Engineering, Norwegian University of Science and Technology
†Department of Electrical Engineering, Stanford University

Contents

1 Introduction

Retail electricity tariffs typically combine an energy charge, proportional to consumption, with a peak power charge (also called a demand or capacity charge) that reflects the customer's contribution to network peaks [1, 2]. These charges have traditionally applied to commercial and industrial customers, but as residential electrification accelerates (heat pumps, electric vehicles, on-site solar), several European countries have extended them to households, with different designs now in use across the Nordic countries, the Netherlands, and parts of central Europe [20, 21].

Norway introduced such a tariff on 1 July 2022, motivated by the need to align charges with network costs and defer grid upgrades [15]; capacity-based residential tariffs in the Norwegian context have been studied in detail, see, *e.g.*, [16]. Unlike a traditional demand charge proportional to the single maximum power draw in a month, the Norwegian tariff applies a *tiered* charge to the *average of the N largest daily peak powers* (typically $N = 3$). The monthly charge is piecewise constant in this average, with jumps at a small number of thresholds. In addition to this grid charge, the energy portion of the bill combines a small time-of-use retail component with the Nord Pool day-ahead wholesale price [19].

A home battery can shift consumption to low-price periods and shave peaks to reduce the grid charge. These two objectives interact: greedy energy arbitrage charges the battery at full rate whenever prices dip, driving the grid draw into the highest peak tier. Minimizing total cost requires jointly handling energy prices and the piecewise-constant peak cost.

1.1 Our contribution

We show that home battery dispatch under the Norwegian tiered peak power tariff admits a tractable optimization formulation. Solved with perfect foresight of loads and prices, it gives a prescient lower bound on the cost of any implementable policy; run on a rolling horizon with forecasts, it gives an implementable model predictive control (MPC) policy. On one year of hourly data from a home in Trondheim, the MPC policy attains a cost within 1.7% of the prescient bound, and close to three times the savings of the best rule-based policy we consider. To the best of our knowledge, the tariff structure we consider, a tiered piecewise-constant charge on the average of the N largest daily peaks, has not been addressed in the prior literature on storage control under peak power charges, which has focused on linear demand charges proportional to a single monthly maximum.

1.2 Related work

Convex energy management. A unified framework for dynamic energy management based on convex optimization appears in [13], in which generators, storage, and loads are modeled as a network of devices with convex objectives and constraints. MPC variants that handle forecast uncertainty through scenario-based robust optimization are developed in [9]. Our formulation follows this modeling approach; the only source of nonconvexity in our problem is the tiered peak power charge, which introduces integer variables. Flexible and thermostatic loads, the focus of much of the residential home energy management (HEMS) literature [7], fit directly into this framework as additional devices.

Battery scheduling under peak power charges. Most prior work on battery scheduling under peak power charges assumes a linear penalty on the single maximum power over the billing period. Rule-based threshold policies, which discharge the battery whenever the load exceeds a target, are a common baseline [6]. MPC methods that jointly optimize peak reduction and energy cost, in deterministic or stochastic formulations, appear in [11, 12, 14]. Dynamic programming approaches that carry the running peak as part of the state are developed in [17], using value-function approximation to manage the curse of dimensionality. The MPC policy we develop can be viewed as an approximate dynamic programming method in which the finite-horizon optimization plays the role of a rolling value function. The Norwegian tariff we consider differs from the structures studied in this line of work in two respects: it uses the average of the N largest daily peaks rather than a single monthly maximum, and the charge is tiered and piecewise constant rather than linear in this average.

1.3 Outline

In §2 we describe the home energy system and the tiered peak power cost. In §3 we introduce a running example based on data from a home in Trondheim, Norway. In §4 we formulate the prescient problem as an MILP and establish a lower bound on achievable cost. In §5 we develop the MPC policy and compare it against the prescient bound and three rule-based baselines. In §6 we describe the baseline-plus-AR forecasting method used inside MPC. Appendix A reports a sensitivity analysis of MPC performance to the planning horizon, the peak-average parameter N, and the forecast method.

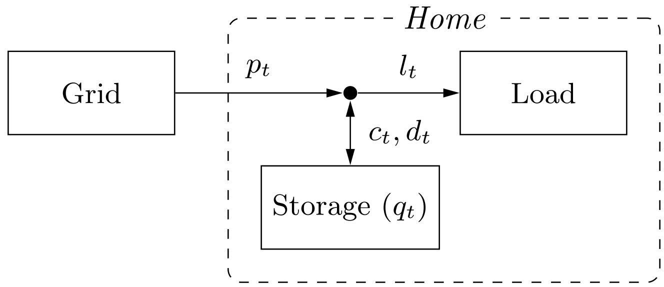

Figure 1: Grid-connected home with storage. The power flows are the grid draw p_t, the charging and discharging powers c_t and d_t, and the net load l_t. The storage device has an internal charge level q_t.

2 Home energy system

We consider a grid-connected home with a battery, as shown in figure 1.

2.1 Power balance

We use hourly time steps, indexed by $t = 1, 2, \ldots, T$. The load l_t is a net value that can include, for example, photovoltaic generation; it is usually nonnegative. The storage power splits into a nonnegative charging rate c_t and a nonnegative discharging rate d_t. The grid power p_t is the power drawn from the grid, with $0 \leq p_t \leq P$, where $P > 0$ is the maximum grid power. Power balance requires

$$p_t + d_t - l_t - c_t = 0, \quad t = 1, \ldots, T-1.$$

2.2 Storage device

The storage device has a charge level q_t, with dynamics

$$q_{t+1} = \eta_s q_t + h(\eta_c c_t - (1/\eta_d) d_t), \quad t = 1, \ldots, T-1,$$

where $\eta_s, \eta_c, \eta_d \in (0, 1)$ are the storing, charging, and discharging efficiencies, and h is the time period in hours. The initial charge level $q_1 = q_{\text{init}}$ is given; we optionally impose a terminal constraint $q_T = q_{\text{final}}$. The charge level satisfies

$$0 \leq q_t \leq Q, \quad t = 1, \ldots, T,$$

where $Q \geq 0$ is the storage capacity, and the charge and discharge rates satisfy

$$0 \leq c_t \leq C, \quad 0 \leq d_t \leq D, \quad t = 1, \ldots, T-1,$$

where C and D are the positive maximum rates.

2.3 Cost function

The cost function consists of an energy charge and a peak power charge.

Energy charge. The energy charge is $h\sum_{t=1}^{T-1}(\lambda_t^{\text{tou}}+\lambda_t^{\text{da}})p_t$, where λ_t^{tou} are time-of-use prices that vary by time of day and season, and λ_t^{da} are day-ahead prices set daily through an auction market.

Peak power charge. The peak power charge is assessed monthly. We assume the time horizon spans K whole months, denoted $k = 1, \ldots, K$. Let m_k denote the vector of daily maximum powers in month k, and let z_k denote the average of the N largest entries of m_k, where N is a parameter (typically $N = 3$). When $N = 1$, z_k is simply the maximum daily peak power over the month. We define the sum-largest function $\psi(u, N)$ as the sum of the largest N components of a vector u, so $z_k = \psi(m_k, N)/N$.

The peak power cost in month k is $\varphi(z_k)$, where φ is piecewise constant with the form

$$\varphi(z) = \begin{cases} \beta_1 & 0 \le z \le r_1 \\ \beta_2 & r_1 < z \le r_2 \\ \vdots & \\ \beta_{L-1} & r_{L-2} < z \le r_{L-1} \\ \beta_L & z > r_{L-1} \end{cases} \tag{1}$$

where $0 < \beta_1 < \beta_2 < \cdots < \beta_L$ are the charges, and $0 < r_1 < r_2 < \cdots < r_{L-1}$ are the thresholds. When $\varphi(z_k) = \beta_j$, we say that the peak power cost is in tier j. We set $r_L = P$, so the last condition in (1) can be expressed as $r_{L-1} < z \le r_L$.

Total cost. The overall cost is

$$h\sum_{t=1}^{T-1}(\lambda_t^{\text{tou}}+\lambda_t^{\text{da}})p_t + \sum_{k=1}^{K}\varphi(z_k).$$

The first term, the energy charge, is linear in the power values, while the second term, the peak power charge, is nonconvex.

2.4 Policies

Information pattern. We assume that the following quantities are known.

- The system parameters P (max grid power), Q (battery capacity), C (max charge rate), D (max discharge rate), and the efficiencies η_s, η_c, and η_d.
- The initial and final charge levels, q_{init} and q_{final}.
- Time-of-use prices $\lambda_1^{\text{tou}}, \dots, \lambda_{T-1}^{\text{tou}}$.
- The parameter N in the peak power calculation.
- The costs $\beta_1, \dots, \beta_L$ and thresholds $r_1, \dots, r_{L-1}$ specifying the peak power cost tiers.

Day-ahead prices λ_t^{da} are announced daily at 13:00 for the following day, so prices are known 12–35 hours in advance. The load l_t is known at the beginning of period t, but future loads are not known.

Prescient policy. All loads $l_1, \dots, l_{T-1}$ and day-ahead prices $\lambda_1^{\text{da}}, \dots, \lambda_{T-1}^{\text{da}}$ are known in advance. The charging and discharging powers are chosen to minimize total cost subject to the constraints above. This policy is not implementable, but provides a lower bound on the cost achievable by any implementable policy.

Implementable policy. An implementable policy respects the information pattern; in period t, it chooses c_t and d_t based on loads and day-ahead prices known at time t. Future values are not known, but can be forecast from past observations.

3 Running example

We illustrate our methods on a running example using real data from a home in Trondheim, Norway, over 2020–2022. We use 2020–2021 data to fit forecast models (see §6) and 2022 data to evaluate policies. The data and source code are available at https://github.com/cvxgrp/home-battery-dispatch.

3.1 The Norwegian electricity tariff

A Norwegian residential electricity bill combines three components: a day-ahead wholesale energy cost, a small time-of-use retail energy cost, and a grid tariff that includes the tiered peak power charge we study here.

Day-ahead prices. Wholesale electricity is traded on the Nord Pool spot market, which clears hourly prices in each bidding zone through a daily auction. Prices for the next day are published at 13:00 for delivery 00:00–23:00. Trondheim is in the NO3 zone. Day-ahead prices show a strong diurnal pattern (morning and evening peaks, mid-day trough) and large seasonal variation; they climbed sharply during the 2021–2022 European energy crisis, as visible in figure 3. In 2022 the day-ahead component was the largest component of the bill, accounting for roughly half of the no-storage cost.

Retail energy charge. Retailers add a small markup to the wholesale price. We use a time-of-use tariff that is lower at night and in winter, representative of residential contracts available in 2022. Variable-price contracts are the most common in Norway; fixed-price and government-capped contracts (the latter introduced in 2022) are a minority. The peak power charge we study applies independently of the energy contract, so our formulation extends directly to these contract types, with the energy arbitrage component trivialized by a constant energy price.

Tiered peak power charge. Norway's tiered residential capacity charge, introduced on 1 July 2022 following a recommendation from the Norwegian energy regulator RME [15], ties the monthly capacity charge to the customer's own peak demand. Each distribution network operator publishes a schedule of tiers, with higher tiers applied when the monthly peak exceeds a threshold. The schedule we use, shown in table 2, is from the distribution network operator for the Trondheim region. The monthly charge depends on z_k, the average of the $N = 3$ largest daily peaks, which smooths out single outlier days but still ties the bill to the customer's contribution to network peaks.

3.2 System parameters

We use one-hour intervals, so $h = 1$. The hourly loads l_t are shown in figure 2, over the full three years (top) and one week in January 2022 (bottom). The system parameters are

$$P = 20\,\text{kW}, \quad Q = 40\,\text{kWh}, \quad C = 20\,\text{kW}, \quad D = 20\,\text{kW},$$

with efficiencies

$$\eta_s = 0.99998, \quad \eta_c = 0.95, \quad \eta_d = 0.95.$$

A complete charge or discharge takes 2 hours at full rate. The storing efficiency η_s corresponds to 1.5% monthly self-discharge, typical for lithium-ion batteries. The

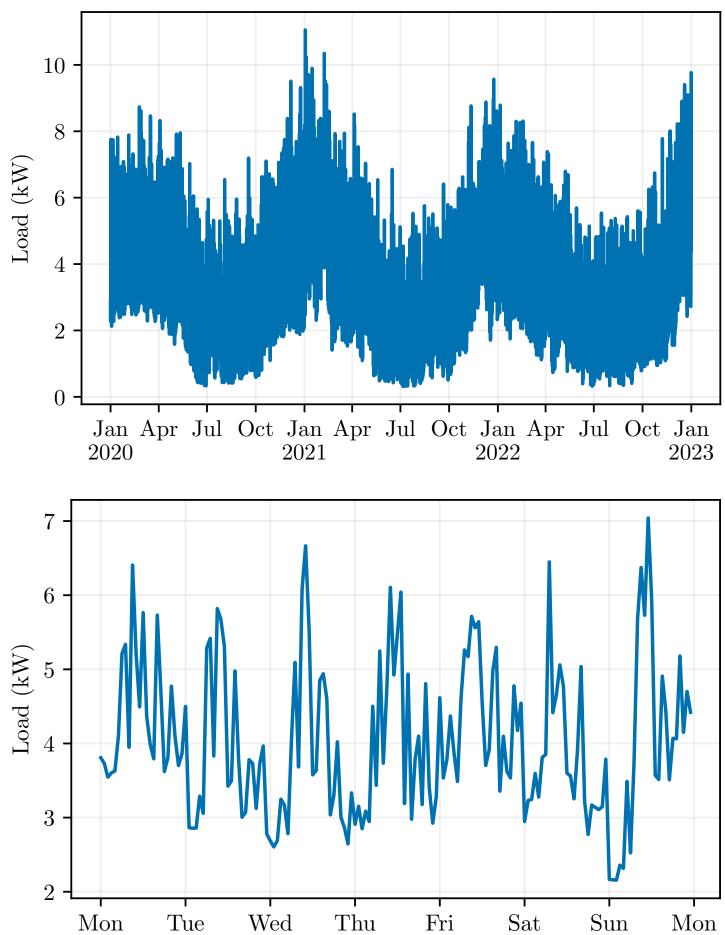

Figure 2: Hourly loads from a home in Trondheim, Norway. *Top.* Three-year period 2020–2022. *Bottom.* One week in January 2022.

Table 1: Time-of-use prices λ_t^{tou} (NOK/kWh).

	Day (hours 6–21)	Night (hours 22–5)
Winter (Jan–Mar)	0.30	0.21
Remaining months	0.39	0.30

Table 2: Tiered peak power cost (NOK/month).

Tier l	Threshold r_l (kW)	Cost β_l (NOK/month)
1	2	83
2	5	147
3	10	252
4	15	371
5	20	490

charging and discharging efficiencies each represent a 5% loss. We set the initial and final charge levels to $q_{\text{init}} = q_{\text{final}} = Q/2$, a symmetric choice. In some experiments we vary Q around its nominal value of 40 kWh to study the effect of storage capacity on cost savings.

3.3 Price parameters

Figure 3 shows day-ahead prices λ_t^{da} for Trondheim, in Norwegian Krone (NOK) per kWh, over the full three-year period (top) and one week in January 2022 (bottom); data are sourced from Nord Pool [19].

Table 1 shows time-of-use prices, which vary by time of day and season.

Table 2 shows the tiered peak power cost. The monthly charge is β_l when $r_{l-1} < z_k \leq r_l$, where z_k is the average of the $N = 3$ largest daily peak powers in month k.

3.4 No storage baseline

Without storage, the cost in 2022 is 25,052 NOK, of which 22,028 NOK is the energy charge and 3,024 NOK is the peak power charge.

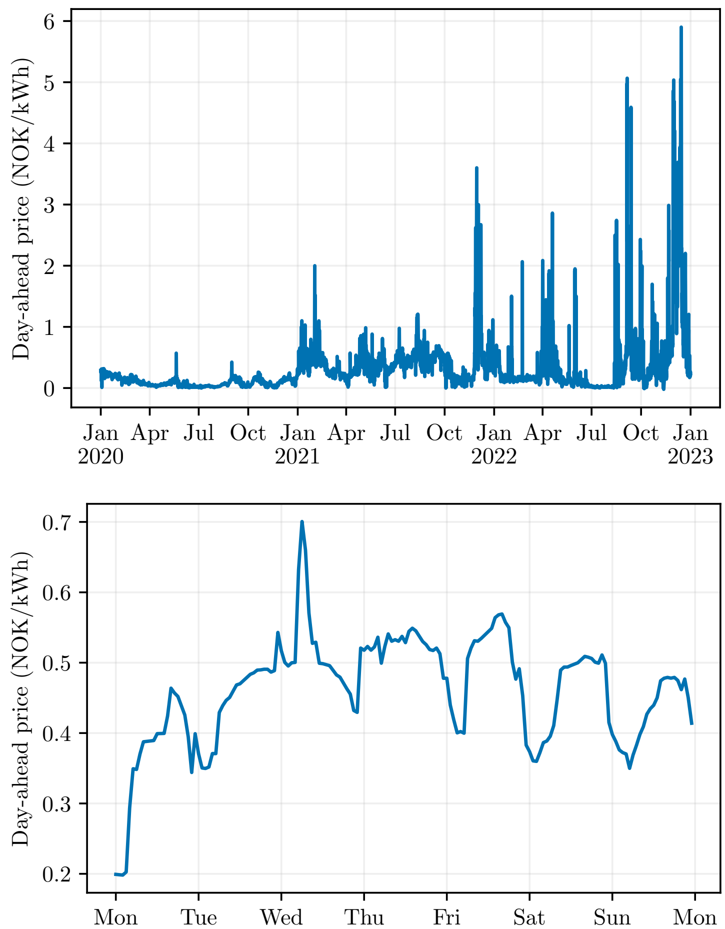

Figure 3: Day-ahead prices λ_t^{da} for Trondheim, Norway. *Top.* Three-year period 2020–2022. *Bottom.* One week in January 2022.

4 Prescient problem

In the prescient problem, we assume all parameters are known, including future loads and day-ahead prices. With $h = 1$ suppressed from the notation, the prescient problem is expressed as the optimization problem

$$
\begin{array}{lll}
\text{minimize} & \sum_{t=1}^{T-1}(\lambda_t^{\text{tou}} + \lambda_t^{\text{da}})p_t + \sum_{k=1}^{K}\varphi(z_k) & \\
\text{subject to} & p_t + d_t - l_t - c_t = 0, & t = 1,\ldots,T-1 \\
& q_{t+1} = \eta_s q_t + \eta_c c_t - (1/\eta_d)d_t, & t = 1,\ldots,T-1 \\
& q_1 = Q/2, \quad q_T = Q/2 & \\
& 0 \le q_t \le Q, & t = 1,\ldots,T \\
& 0 \le p_t \le P, & t = 1,\ldots,T-1 \\
& 0 \le c_t \le C, & t = 1,\ldots,T-1 \\
& 0 \le d_t \le D, & t = 1,\ldots,T-1
\end{array}
\tag{2}
$$

with variables $p, c, d \in \mathbf{R}^{T-1}$ and $q \in \mathbf{R}^T$, giving about $4T$ scalar variables. Here $z_k = \psi(m_k, N)/N$, where m_k is the vector of daily maxima of p_t in month k, and φ is the piecewise constant peak power cost function (1).

4.1 MILP formulation

The piecewise constant function φ can be modeled by introducing binary variables $s_{lk} \in \{0, 1\}$ for tier $l = 1, \ldots, L$ and month $k = 1, \ldots, K$, with

$$
\varphi(z_k) = \sum_{l=1}^{L} \beta_l s_{lk}, \qquad z_k \le \sum_{l=1}^{L} r_l s_{lk}, \qquad \sum_{l=1}^{L} s_{lk} = 1.
$$

When $s_{lk} = 1$, z_k is assigned to tier l with cost β_l. This adds LK binary variables.

The resulting problem is a mixed-integer convex problem (MICP): all constraints are convex except the integrality requirement $s_{lk} \in \{0, 1\}$. The objective is linear, most constraints are linear equalities or inequalities, and the only nonlinearities involve m_k (each entry is the maximum of a subset of the p_τ, hence convex) and $z_k = \psi(m_k, N)/N$, which is convex because $\psi(\cdot, N)$ is convex and nondecreasing [4, §3.2.3].

The max and sum-largest functions are piecewise linear and can be represented via linear inequalities, so (2) can be transformed to an equivalent MILP. Domain-specific languages (DSLs) for convex optimization such as CVXPY [8] automate these transformations and generate an MILP that can be passed to solvers such as Gurobi [18] or open-source alternatives.

4.2 CVXPY implementation

We show how to formulate and solve problem (2) using CVXPY. We assume the parameters from §3 and data vectors `l`, `tou_prices`, `da_prices`, and datetime index `dt` are defined.

```
1 import cvxpy as cp
2 from pandas import unique
3
4 p = cp.Variable(T-1, nonneg=True)
5 c = cp.Variable(T-1, nonneg=True)
6 d = cp.Variable(T-1, nonneg=True)
7 q = cp.Variable(T, nonneg=True)
8 s = cp.Variable((K, L), boolean=True)
9
10 cons = [p + d - l - c == 0,
11         q[1:] == eta_s*q[:-1] + eta_c*c - d/eta_d,
12         q[0] == Q/2, q[-1] == Q/2,
13         q <= Q, p <= P, c <= C, d <= D]
14
15 energy_cost = cp.sum(cp.multiply(tou_prices + da_prices, p))
16
17 peak_cost = 0
18 for k, month in enumerate(sorted(dt.month.unique())):
19     days = unique(dt[dt.month == month].date)
20     m_k = [cp.max(p[dt.date == day]) for day in days]
21     z_k = cp.sum_largest(cp.hstack(m_k), N) / N
22     peak_cost += beta @ s[k]
23     cons += [z_k <= r @ s[k], cp.sum(s[k]) == 1]
24
25 prob = cp.Problem(cp.Minimize(energy_cost + peak_cost), cons)
26 prob.solve()
```

Lines 4–8 define variables matching the dimensions in (2). Lines 10–13 encode constraints on power balance, storage dynamics, and bounds. The loop (lines 18–23) builds the tiered peak power cost for each month and adds the corresponding tier constraints, computing daily maxima m_k and the average of the N largest via `cp.sum_largest`.

In roughly 20 lines we formulate and solve problem (2). CVXPY automatically transforms expressions like `cp.max` and `cp.sum_largest` into equivalent linear constraints before passing the problem to an MILP solver.

Table 3: Annual electricity costs (NOK) for 2022 with 40 kWh storage.

	Energy	Peak power	Total	Savings
No storage	22,028	3,024	25,052	—
Prescient	19,399	1,805	21,204	15.4%

4.3 Running example

We solve problem (2) using 2022 data with CVXPY and Gurobi. After compilation, the problem has 35,783 continuous variables and 60 binary variables. On a laptop with Apple M1 Pro processor, it solves in about 6 seconds.

Costs. Table 3 compares costs with and without storage. The prescient cost is 21,204 NOK, a savings of 15.4% over the no-storage baseline. Roughly two-thirds of the savings come from reduced energy charges.

Power flows. Figure 4 shows the optimal power flows over the year. The policy achieves tier 2 in ten months, tier 1 in July when loads are lowest, and tier 3 only in December when winter peaks are highest. Figure 5 shows one week in detail; the battery charges when prices are low and discharges during load peaks, reducing both energy and peak power costs.

Savings versus capacity. Figure 6 shows annual savings versus storage capacity for 2020–2022. The curves exhibit diminishing returns; in 2022, doubling capacity from 20 to 40 kWh increases savings from 12% to 15%. Without the binary tier constraints, the problem is an LP in Q, so the curves are piecewise-linear and concave (a standard LP sensitivity result); with them, tier transitions introduce small discontinuities but leave the overall shape nearly concave.

Economic considerations. Figure 6 lets us translate operational savings into a break-even battery cost. Averaging across 2020–2022, a 40 kWh battery saves about 69 NOK/kWh per year, which over a ten-year service life implies a simple-payback break-even installed cost of roughly 700 NOK/kWh. Residential lithium-ion systems in Norway are installed at around 8,000–12,000 NOK/kWh before subsidy, so a dedicated home battery does not yet pay back on operational savings alone under this tariff.

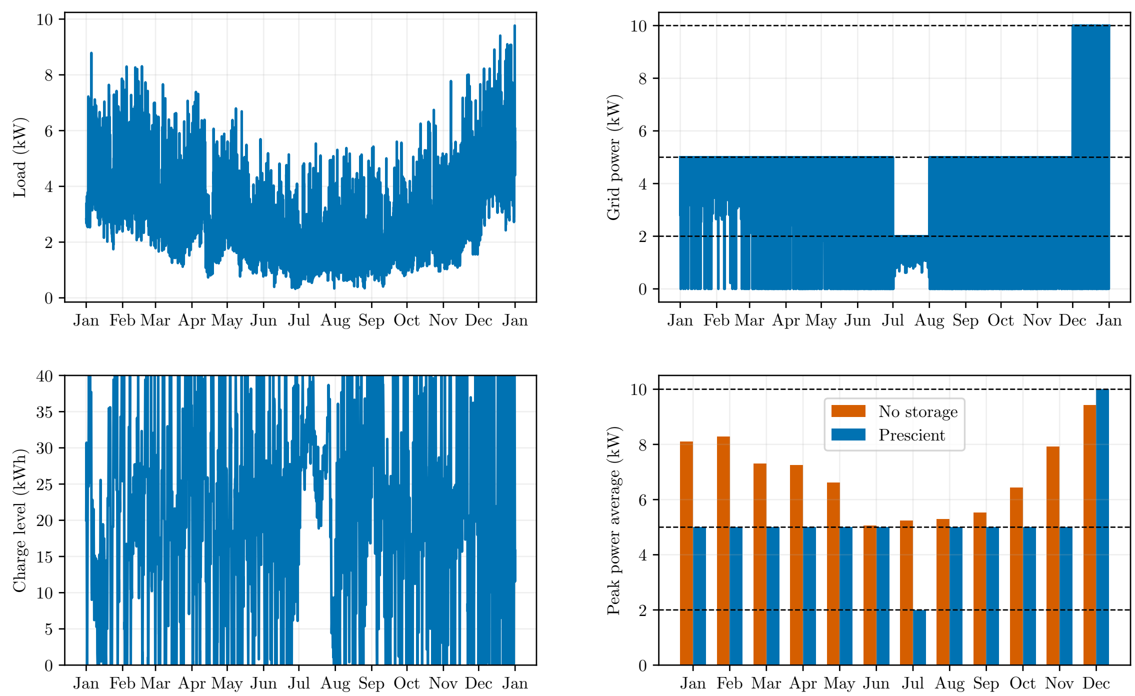

Figure 4: Prescient policy with 40 kWh storage over 2022. Tier thresholds shown as dashed lines. *Top left.* Load. *Top right.* Grid power. *Bottom left.* Charge level. *Bottom right.* Average peak power z_k.

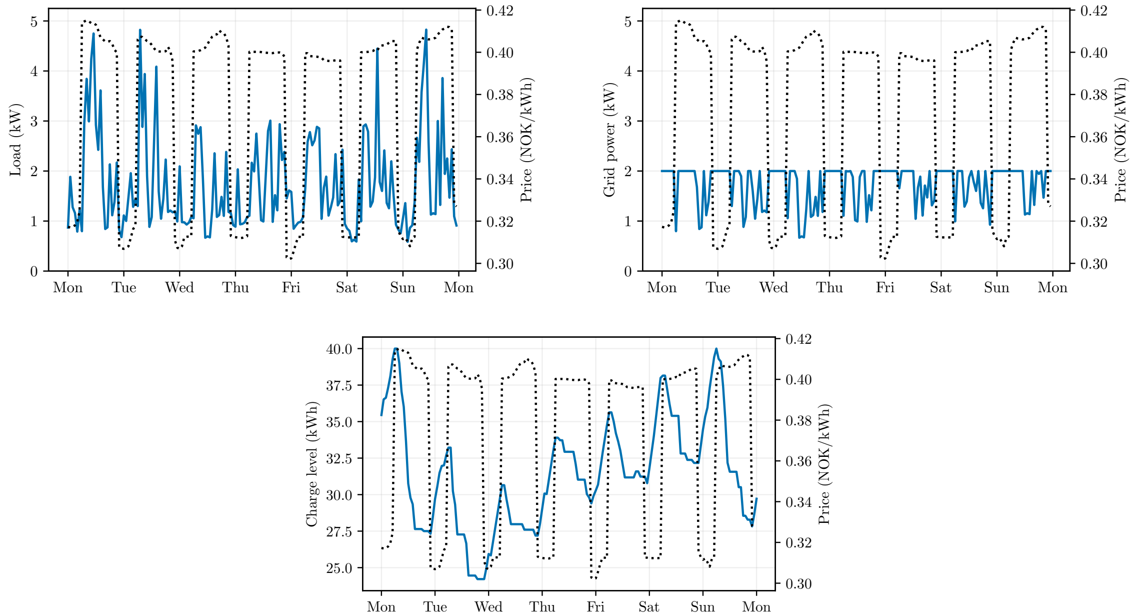

Figure 5: Prescient policy over one week in July 2022. Electricity prices $\lambda_t^{\text{tou}} + \lambda_t^{\text{da}}$ shown as dotted lines. *Top left.* Load. *Top right.* Grid power. *Bottom.* Charge level.

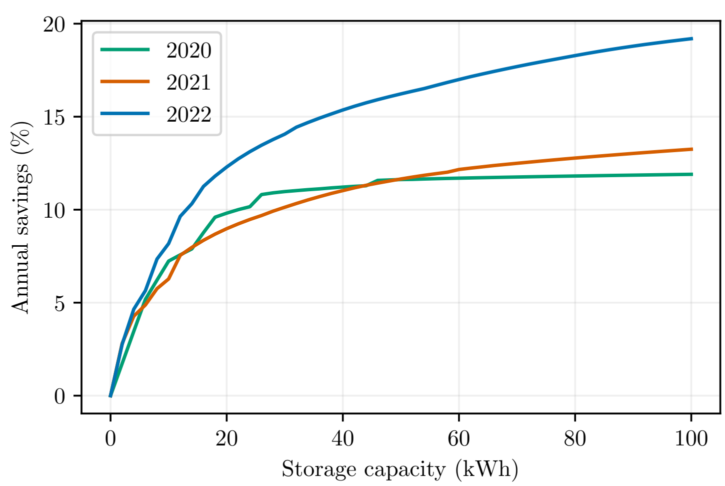

Figure 6: Annual savings versus storage capacity for 2020–2022.

Two observations change this conclusion. Many Norwegian homes already own a 40–80 kWh battery in the form of an electric vehicle; vehicle-to-grid operation using the same controller turns the savings reported here into net savings, with no incremental capital cost [5]. Separately, installed battery costs have declined around 10% per year for a decade, and residential peak tariffs are expected to tighten as electrification accelerates; both trends move break-even in the same direction.

5 Model predictive control

The prescient problem provides a performance bound but requires perfect foresight. We develop an implementable policy using model predictive control (MPC). At each time step, we solve an optimization problem over a finite horizon using forecasts, apply the first action, and repeat with updated information.

5.1 MPC formulation

At hour t, we observe the charge level q_t and form forecasts of loads and day-ahead prices over a planning horizon of H hours. We let $\hat{l}_{\tau|t}$ and $\hat{\lambda}^{\text{da}}_{\tau|t}$ denote the forecasts for period τ made at time t; for $\tau = t$ these are the known current values. The forecasting method is described in §6. Let k denote the current month and M the number of months in the horizon. The MPC problem is

$$\begin{array}{lll}
\text{minimize} & \sum_{\tau=t}^{t+H-1}(\lambda^{\text{tou}}_{\tau} + \hat{\lambda}^{\text{da}}_{\tau|t})p_\tau + \sum_{j=k}^{k+M-1}\varphi(z_{j|t}) & \\
\text{subject to} & p_\tau + d_\tau - \hat{l}_{\tau|t} - c_\tau = 0, & \tau = t,\dots,t+H-1 \\
& q_{\tau+1} = \eta_s q_\tau + \eta_c c_\tau - (1/\eta_d)d_\tau, & \tau = t,\dots,t+H-1 \\
& q_{t+H} = Q/2 & \\
& 0 \le q_\tau \le Q, & \tau = t+1,\dots,t+H \\
& 0 \le p_\tau \le P, & \tau = t,\dots,t+H-1 \\
& 0 \le c_\tau \le C, & \tau = t,\dots,t+H-1 \\
& 0 \le d_\tau \le D, & \tau = t,\dots,t+H-1
\end{array} \tag{3}$$

with variables $p, c, d \in \mathbf{R}^H$ and $q \in \mathbf{R}^{H+1}$, where the initial state q_t is fixed at the observed value. The terminal constraint $q_{t+H} = Q/2$ prevents myopic depletion of the battery.

The peak power terms require care because the monthly peak averages depend on both past realized and future planned powers. Let $m_{j|t}$ be the vector of daily maxima for month j as evaluated at time t. For days completed before t, the entry is the realized daily maximum (a constant). For each day that intersects the horizon

$[t, t+H-1]$, the entry is the maximum of p_τ over the hours in that day, which may combine past realized p_τ with decision variables p_τ for hours $\tau \geq t$. The peak average $z_{j|t}$ is the average of the N largest entries of $m_{j|t}$, or all entries if fewer than N days are represented.

5.2 MILP formulation

Problem (3) has the same structure as (2), with about $4H$ continuous variables. The piecewise constant φ is modeled using binary variables as in §4. A 30-day horizon typically spans two months, giving $2L$ binary variables.

5.3 LP enumeration

An alternative to MILP is enumeration. Since each month selects exactly one tier, we can solve (3) by enumerating over tier assignments. For each assignment the tier variables are fixed and (3) reduces to a linear program. With two months this gives L^2 linear programs; we solve them all and select the minimum-cost solution. This is fast since L is small, and is useful when an MILP solver is unavailable.

5.4 Running example

Simple policies. We compare MPC against three rule-based policies that use only current observations, without forecasts or optimization. All three are clipped to respect charge-level and rate limits.

The *peak shaving* policy discharges when the load exceeds a seasonal target κ_m and charges otherwise,

$$(d_t, c_t) = \begin{cases} (\min\{D,\, l_t - \kappa_m\},\ 0) & l_t > \kappa_m, \\ (0,\ \min\{C,\, \kappa_m - l_t\}) & l_t \leq \kappa_m, \end{cases}$$

where κ_m depends on the current month. We use $\kappa_m = 10\,\text{kW}$ in winter (December–February), $\kappa_m = 5\,\text{kW}$ in shoulder months (March–May, September–November), and $\kappa_m = 2\,\text{kW}$ in summer (June–August), matching the upper boundary of tier 3, tier 2, and tier 1, respectively. This captures the fact that loads are much higher in Norwegian winters than in summers.

The *energy arbitrage* policy uses the current price to decide when to charge and discharge. Let $\lambda_t = \lambda_t^{\text{tou}} + \lambda_t^{\text{da}}$ be the retail price at hour t, and $\bar{\lambda}_t$ its median over the calendar day that contains t. The policy charges at full rate when the current price is below the median and discharges against the load otherwise,

$$(d_t, c_t) = \begin{cases} (0,\ \min\{C,\ (Q - q_t)/\eta_c\}) & \lambda_t < \bar{\lambda}_t, \\ (\min\{D,\ l_t\},\ 0) & \lambda_t \geq \bar{\lambda}_t. \end{cases}$$

The threshold is the daily median rather than a fixed hour schedule, so the policy tracks the prevailing day-ahead price curve.

The *capped arbitrage* policy merges the two by capping the grid draw at the seasonal target κ_m during charging,

$$(d_t, c_t) = \begin{cases} (0,\ \min\{C,\ (Q - q_t)/\eta_c,\ \kappa_m - l_t\}) & \lambda_t < \bar{\lambda}_t, \\ (\min\{D,\ l_t\},\ 0) & \lambda_t \geq \bar{\lambda}_t. \end{cases}$$

Charging is thus clipped both to the charge level and to the tier cap, so the policy exploits energy arbitrage when prices are low but never drives the grid draw above its seasonal tier.

MPC parameters. We run MPC with horizon $H = 720$ hours (30 days) and $N = 3$, matching the tariff definition, using the forecasts described in §6. With $L = 5$ tiers, the horizon spans two months, giving $2L = 10$ binary variables. Solving the MILP with Gurobi, each step takes around 0.2 seconds. Sensitivity to H, N, and the forecast method is explored in Appendix A.

Costs. Table 4 compares all policies. MPC achieves a total cost of 21,568 NOK, a savings of 13.9% relative to no storage and within 1.7% of the prescient bound. Peak shaving achieves 3.3% savings by reducing peak power charges in shoulder and summer months but forgoes energy arbitrage; even with a seasonal target, the policy is sometimes too aggressive in summer (the battery depletes serving routine load) and too passive in winter (loads rarely exceed the target, so the battery stays idle). Energy arbitrage reduces energy cost by nearly a thousand NOK but, by charging at full rate whenever prices fall below the daily median, drives peak power into tier 5 every month, giving a total cost *higher* than no storage. Capped arbitrage, which clips the grid draw during low-price charging to the seasonal target, avoids this failure and achieves 5.0% savings.

The gap between the best rule-based policy (5.0%) and MPC (13.9%) is almost a factor of three, and quantifies the value of forecasting and joint optimization. To isolate the contribution of modeling the peak power charge, we also report an MPC

Table 4: Comparison of policies for 2022 with 40 kWh storage. Costs in NOK.

	Energy	Peak power	Total	Savings
No storage	22,028	3,024	25,052	—
Peak shaving	22,009	2,225	24,234	3.3%
Energy arbitrage	21,088	5,880	26,968	−7.7%
Capped arbitrage	20,779	3,024	23,803	5.0%
MPC (energy-only)	18,774	5,880	24,654	1.6%
MPC	19,279	2,289	21,568	13.9%
Prescient	19,399	1,805	21,204	15.4%

variant that ignores it (energy-only). This policy achieves the lowest energy cost but, like energy arbitrage, hits tier 5 every month, yielding savings of only 1.6%. The jump from 1.6% to 13.9% shows that most of the value of MPC comes from jointly handling energy and peak charges; arbitrage alone is not enough.

Power flows. Figure 7 shows the MPC power flows over the year. MPC achieves tier 2 in fewer months than prescient, since forecast errors make it harder to anticipate the highest load days. The weekly detail in figure 8 shows similar coordination to prescient, with timing differences due to uncertainty in load and price forecasts.

6 Forecasting

In this section we describe a simple method to forecast a scalar time series $x_1, x_2, \ldots$ using historical data, following [13, Appendix A].

6.1 The baseline-residual forecast

The baseline-residual forecast is given by

$$\hat{x}_{\tau|t} = b_\tau + \hat{r}_{\tau|t},$$

where $\hat{x}_{\tau|t} \in \mathbf{R}$ is the prediction of quantity x at time $\tau \geq t$, made at time t. The forecast is the sum of a seasonal baseline b_τ, capturing periodic patterns (diurnal, weekly, annual), and an autoregressive residual $\hat{r}_{\tau|t}$ that accounts for short-term deviations from the baseline. The baseline depends only on τ; the residual depends on both τ and the time t at which the forecast is made.

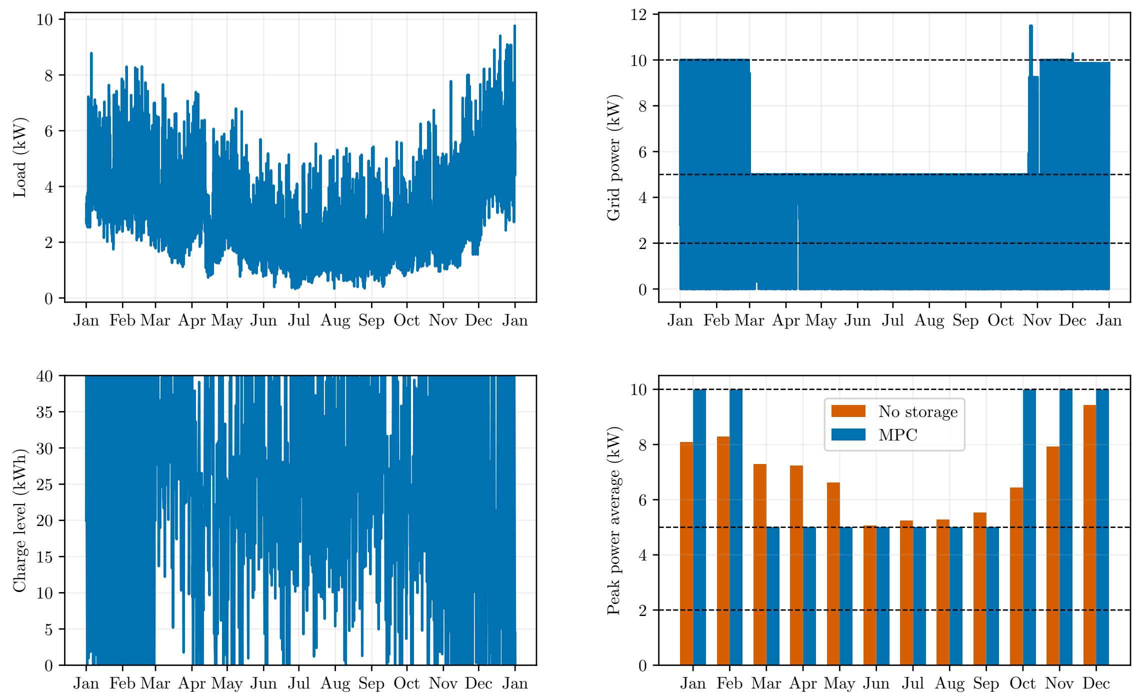

Figure 7: MPC policy with 40 kWh storage over 2022. Tier thresholds shown as dashed lines. *Top left.* Load. *Top right.* Grid power. *Bottom left.* Charge level. *Bottom right.* Average peak power z_k.

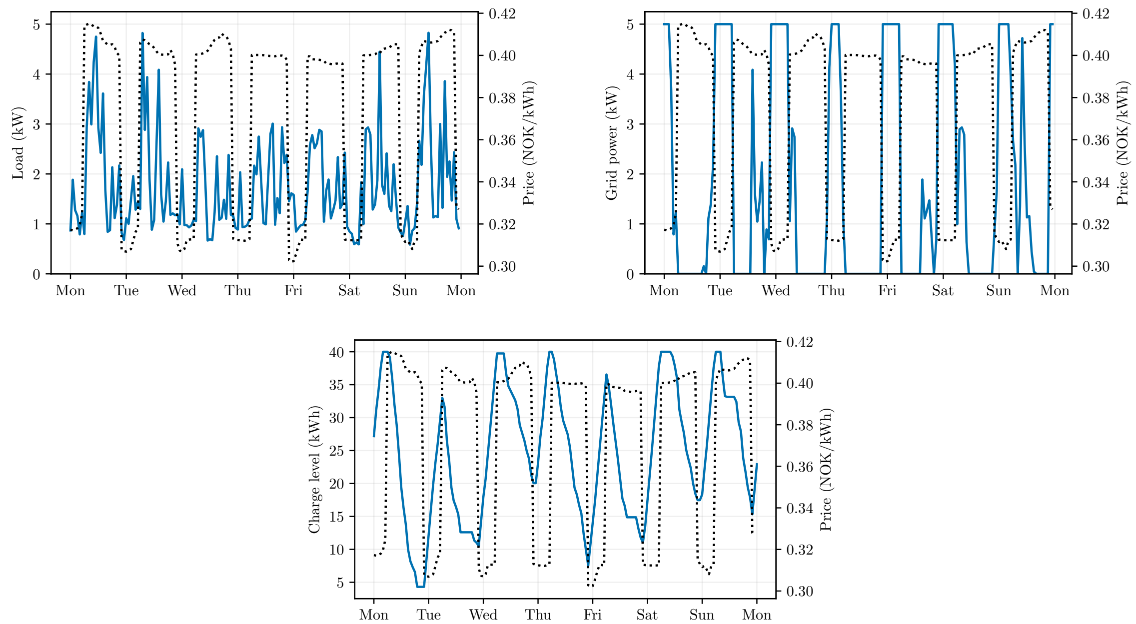

Figure 8: MPC policy over one week in July 2022. Electricity prices $\lambda_t^{\text{tou}} + \lambda_t^{\text{da}}$ shown as dotted lines. *Top left.* Load. *Top right.* Grid power. *Bottom.* Charge level.

Baseline component. A simple model for the baseline is a sum of K sinusoids,

$$b_t = \beta_0 + \sum_{k=1}^{K} \left(\alpha_k \sin\left(\frac{2\pi t}{P_k}\right) + \beta_k \cos\left(\frac{2\pi t}{P_k}\right) \right),$$

where α_k and β_k are the coefficients and P_k are the periods. In the usual case of Fourier series, the periods have the form $P_k = P/k$, where P is the fundamental period. Here we include terms for daily, weekly, and annual variation.

To fit the $2K+1$ coefficients $\beta_0, \alpha_1, \beta_1, \ldots, \alpha_K, \beta_K$ to historical training data $x_1, \ldots, x_T$, we use a pinball (quantile) loss with ℓ_2 regularization. The pinball loss for quantile $\eta \in [0,1]$ is

$$L_\eta(u) = \max\{\eta u, (\eta - 1)u\} = (\eta - 1/2)u + (1/2)|u|.$$

Being piecewise linear, it is less sensitive to outliers than squared error. Choosing $\eta < 0.5$ biases forecasts toward overestimation, while $\eta > 0.5$ biases toward underestimation. We find the coefficients by minimizing

$$\sum_{t=1}^{T} L_\eta(b_t - x_t) + \lambda \sum_{k=1}^{K} \nu_k(\alpha_k^2 + \beta_k^2),$$

where $\lambda > 0$ is the regularization parameter and $\nu_k = k^2$ penalizes higher harmonics more heavily. The ℓ_2 penalty shrinks coefficients without zeroing them out, since all harmonics contribute to a smooth baseline. The problem is convex and readily solved; good values for η and λ can be chosen by cross-validation [3, §7.10], [10, §13.2].

Residual component. The residuals are $r_t = x_t - b_t$, $t = 1, \ldots, T$. We fit an autoregressive (AR) model to predict the next L residuals from the previous M,

$$(\hat{r}_{t+1|t}, \ldots, \hat{r}_{t+L|t}) = \Gamma(r_{t-M+1}, \ldots, r_t),$$

where $\Gamma \in \mathbf{R}^{L \times M}$ is the parameter matrix. We find Γ by minimizing

$$\sum_{t=1}^{T} \sum_{\tau=t+1}^{t+L} L_\eta(\hat{r}_{\tau|t} - r_\tau) + \lambda \|\Gamma\|_F^2,$$

where $\|\cdot\|_F$ is the Frobenius norm; good values for η and λ can be chosen by cross-validation.

6.2 Running example

Load forecasting. We apply the baseline-residual method to load. The baseline captures diurnal (24h), weekly, and annual periodicities with 4 harmonics each, giving periods

$$\begin{aligned}
&P_1 = 24/1,\ P_2 = 24/2,\ P_3 = 24/3,\ P_4 = 24/4,\\
&P_5 = 168/1,\ P_6 = 168/2,\ P_7 = 168/3,\ P_8 = 168/4,\\
&P_9 = 8760/1,\ P_{10} = 8760/2,\ P_{11} = 8760/3,\ P_{12} = 8760/4.
\end{aligned}$$

We fit these 25 parameters on two years of hourly data (2020–2021). Figure 9 shows the baseline component against actual load for one week each in January and June 2022.

For the residual component, we fit an AR model to predict residuals over the next 23 hours from the previous 24. The 23×24 parameter matrix Γ is fit using quantile regression with ℓ_2 regularization on the same training period. Figure 10 compares the baseline component alone to the full forecast (baseline + AR) for a test day in May 2022.

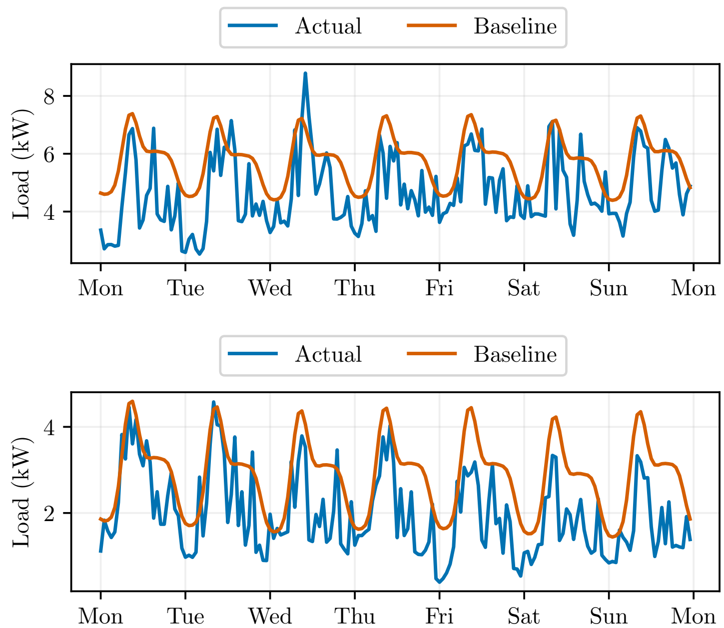

Figure 9: Baseline component (orange) and actual load (blue). *Top.* One week in January 2022. *Bottom.* One week in June 2022.

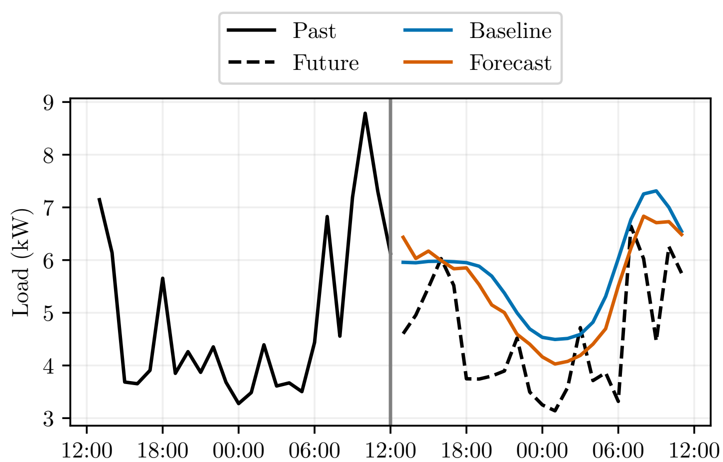

Figure 10: Load forecasts for a test day in May 2022, with vertical line marking the forecast hour. Black solid and dashed lines show realized and future load; blue shows the baseline component; orange the forecast (baseline + AR).

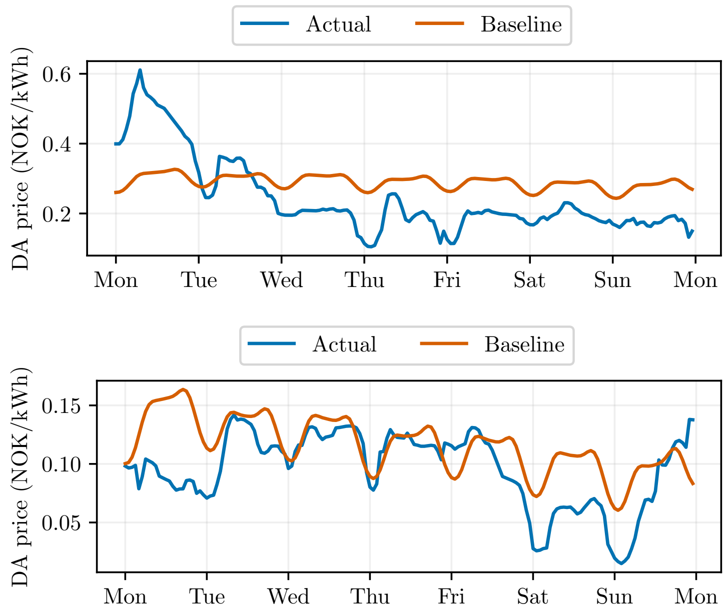

Figure 11: Baseline component (orange) and actual day-ahead prices (blue). *Top.* One week in January 2022. *Bottom.* One week in June 2022.

Forecasting day-ahead prices. We apply the same method to day-ahead prices. Since prices are announced at 13:00 for the next day, the known horizon varies from 12 to 35 hours depending on the time of day. Figure 11 shows the baseline component against actual prices for one week each in January and June 2022. Figure 12 compares the baseline component alone to the full forecast (baseline + AR) for a test day in May 2022.

7 Conclusions

We have developed an MPC policy for managing battery storage in a grid-connected home under an electricity tariff that includes energy charges and a tiered peak power charge. The peak charge is based on the average of the N largest daily peaks in each month. This tiered structure leads to a mixed-integer linear program that can be solved directly, or by enumerating tier assignments and solving a sequence of linear programs. Both approaches are fast and reliable. Numerical experiments on real data from a home in Trondheim, Norway, with a 40 kWh battery, show that MPC achieves annual savings of 14% compared to a no-storage baseline, within 1.7% of

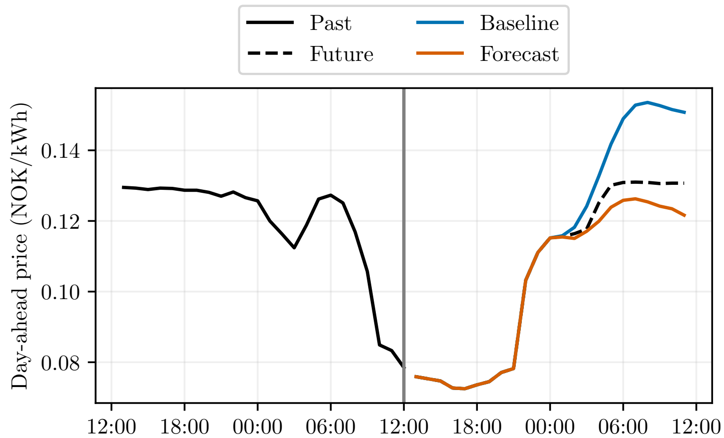

Figure 12: Price forecasts for a test day in May 2022, with vertical line marking the forecast hour. Black solid and dashed lines show realized and future prices; blue shows the baseline component; orange the forecast (baseline + AR).

the prescient bound, and nearly three times the savings of the best rule-based policy we consider.

The methodology extends to several settings beyond the specific running example. The same MILP formulation applies to any tariff defined by piecewise-constant peak charges and to energy assets other than chemical batteries, such as thermal storage and electric vehicle batteries under V2G. Evaluating the impact of battery degradation models and exploring stochastic variants that explicitly represent forecast uncertainty are natural directions for follow-on work.

Acknowledgments

David Pérez-Piñeiro was supported by the Research Council of Norway and HighEFF, an eight-year Research Centre operating under the FME scheme (Centre for Environment-friendly Energy Research, Grant No. 257632). Stephen Boyd was partially supported by ACCESS (AI Chip Center for Emerging Smart Systems), sponsored by InnoHK funding, Hong Kong SAR, and by Office of Naval Research grant N00014-22-1-2121.

A Sensitivity analysis

We study the sensitivity of MPC performance to three design choices: the planning horizon H, the peak power parameter N, and the forecast method. All other parameters are held at their default values from §5.

MPC horizon. Table 5 shows MPC cost as a function of the planning horizon H. A horizon of one day ($H = 24$) performs poorly because the controller cannot anticipate peak power charges beyond the current day; the peak power cost is 3,633 NOK at $H = 24$, versus 2,289 at $H = 720$. Performance improves with longer horizons up to $H = 720$ (30 days), which spans a full billing month. Doubling the horizon to $H = 1440$ does not improve performance further; the natural planning horizon is one billing month, since peak power charges reset monthly.

Table 5: MPC cost versus planning horizon H. Costs in NOK.

H (hours)	Energy	Peak power	Total	Savings
24 (1 day)	19,213	3,633	22,846	8.8%
168 (1 week)	19,234	2,618	21,852	12.8%
360 (2 weeks)	19,280	2,513	21,793	13.0%
720 (30 days)	19,279	2,289	21,568	13.9%
1440 (60 days)	19,352	2,289	21,641	13.6%

Peak power parameter N. The tariff defines the monthly peak average z_k as the mean of the N largest daily maxima. Table 6 compares $N = 1$ and $N = 3$. The difference is negligible (5 NOK over the year). With $N = 1$, the controller optimizes against the single worst day in each month; with $N = 3$ (the actual tariff definition), it can tolerate occasional peaks as long as the average remains low. The robustness follows because limiting the daily maximum also limits the average of the top three.

Table 6: MPC cost versus peak power parameter N. Costs in NOK.

N	Energy	Peak power	Total	Savings
1	19,274	2,289	21,563	13.9%
3	19,279	2,289	21,568	13.9%

Forecast method. Table 7 compares MPC using the full baseline+AR forecast to a baseline-only variant that omits the autoregressive residual component. The improvement from the AR correction is modest (90 NOK, or 0.4%) and comes entirely from reduced peak power charges (2,289 versus 2,394). The AR model helps the controller anticipate near-term load peaks, which improves peak shaving; the seasonal baseline alone captures the price patterns that drive energy arbitrage.

Table 7: MPC cost versus forecast method. Costs in NOK.

Forecast	Energy	Peak power	Total	Savings
Baseline only	19,264	2,394	21,658	13.5%
Baseline + AR	19,279	2,289	21,568	13.9%

References

[1] S. Berg and A. Savvides, “The theory of maximum kW demand charges for electricity,” *Energy Economics*, vol. 5, no. 4, pp. 258–266, 1983.

[2] S. Henderson, “The economics of electricity demand charges,” *The Energy Journal*, vol. 4, no. Special Issue, 1983.

[3] T. Hastie, R. Tibshirani, and J. Friedman, *The elements of statistical learning: Data mining, inference, and prediction*. Springer, 2001.

[4] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.

[5] W. Kempton and J. Tomić, “Vehicle-to-grid power fundamentals: Calculating capacity and net revenue,” *Journal of Power Sources*, vol. 144, no. 1, pp. 268–279, 2005.

[6] G. Karmiris and T. Tengnér, “Peak shaving control method for energy storage,” ABB Corporate Research Center, Västerås, Sweden, Tech. Rep., 2013.

[7] M. Beaudin and H. Zareipour, “Home energy management systems: A review of modelling and complexity,” *Renewable and Sustainable Energy Reviews*, vol. 45, pp. 318–335, 2015.

[8] S. Diamond and S. Boyd, “CVXPY: A Python-embedded modeling language for convex optimization,” *Journal of Machine Learning Research*, vol. 17, no. 83, pp. 1–5, 2016.

[9] M. Wytock, N. Moehle, and S. Boyd, “Dynamic energy management with scenario-based robust MPC,” in *2017 American Control Conference (ACC)*, IEEE, 2017, pp. 2042–2047.

[10] S. Boyd and L. Vandenberghe, *Introduction to applied linear algebra: Vectors, matrices, and least squares*. Cambridge University Press, 2018.

[11] R. Kumar, M. Wenzel, M. Ellis, M. ElBsat, K. Drees, and V. Zavala, “A stochastic model predictive control framework for stationary battery systems,” *IEEE Transactions on Power Systems*, vol. 33, no. 4, pp. 4397–4406, 2018.

[12] M. E. Raoufat, B. Asghari, and R. Sharma, “Model predictive BESS control for demand charge management and PV-utilization improvement,” in *2018 IEEE Power & Energy Society Innovative Smart Grid Technologies Conference (ISGT)*, IEEE, 2018, pp. 1–5.

[13] N. Moehle, E. Busseti, S. Boyd, and M. Wytock, "Dynamic energy management," *Large Scale Optimization in Supply Chains and Smart Manufacturing: Theory and Applications*, pp. 69–126, 2019.

[14] M. Risbeck and J. Rawlings, "Economic model predictive control for time-varying cost and peak demand charge optimization," *IEEE Transactions on Automatic Control*, vol. 65, no. 7, pp. 2957–2968, 2020.

[15] K. R. Verlo, B. A. Fladen, A. Meling, and U. Sira, "Oppsummering av høring og anbefaling til endringer i nettleiestrukturen," Reguleringsmyndigheten for energi (RME), Norges vassdrags- og energidirektorat (NVE), Tech. Rep. Rapport 6/2020, 2020.

[16] S. Bjarghov, H. Farahmand, and G. Doorman, "Capacity subscription grid tariff efficiency and the impact of uncertainty on the subscribed level," *Energy Policy*, vol. 165, p. 112 972, 2022.

[17] M. E. Cholette, L. Liu, B. Jeddi, and Y. Mishra, "Battery dispatching for end users with on-site renewables and peak demand charges—An approximate dynamic programming approach," *IEEE Transactions on Control Systems Technology*, vol. 30, no. 5, pp. 2100–2114, 2022.

[18] Gurobi Optimization, LLC, *Gurobi Optimizer Reference Manual*, 2023.

[19] Nord Pool, *Nord Pool: Power market operations*, 2023.

[20] K. Wang, X. Lai, F. Wen, P. P. Singh, S. Mishra, and I. Palu, "Dynamic network tariffs: Current practices, key issues and challenges," *Energy Conversion and Economics*, vol. 4, no. 1, pp. 23–35, 2023.

[21] M. Hofmann, S. Bjarghov, H. Sæle, and K. B. Lindberg, "Grid tariff design and peak demand shaving: A comparative tariff analysis with simulated demand response," *Energy Policy*, vol. 198, p. 114 475, 2025.