thermo

MDPI

Article

Forecast-Aware Dispatch and Thermodynamic Assessment of Storage-Enabled Diesel Generator Displacement in Grid-Constrained Residential Systems

Ahmed Mohammed [1,*], Qian Zhang [1], Badr Abdullah [2], Omar Aldhaibani [3], Jeff Cullen [2], Alessandro Raschella [3] and Noor Abdalghani [4]

[1] School of Engineering, Liverpool John Moores University, Liverpool L3 3AF, UK; q.zhang@ljmu.ac.uk
[2] School of Civil Engineering and Built Environment, Liverpool John Moores University, Liverpool L3 3AF, UK; b.m.abdullah@ljmu.ac.uk (B.A.); j.d.cullen@ljmu.ac.uk (J.C.)
[3] School of Computer Science and Mathematics, Liverpool John Moores University, Liverpool L3 3AF, UK; o.a.aldhaibani@ljmu.ac.uk (O.A.); a.raschella@ljmu.ac.uk (A.R.)
[4] Department of Computer Science, Tikrit University, Tikrit 34001, Iraq; noor.a.abdalghane35534@st.tu.edu.iq
* Correspondence: a.s.mohammed@ljmu.ac.uk

Abstract

Diesel generators that bridge scheduled grid outages in Iraqi households run mostly at low part load, where their electrical efficiency is about 12%. This study treats a 10 kW generator as a thermal energy converter and develops a forecast-aware controller for a household battery, using one year of measured load and PV data and a Baghdad-type outage timetable. Every hour the controller forms three load scenarios from a 24 h forecast and the out-of-fold quantiles of its error, and it solves a dynamic programme that contains the generator part-load fuel curve. It is compared with no storage, rule-based storage, grid pre-charge, a point-forecast controller and a perfect-foresight bound. A 5 kWh battery reduces diesel use from about 2290 to 212 L/yr (91% against no storage, 82% against rule-based storage), and the efficiency of the remaining generator operation rises from about 12% to 25%. The scenario controller uses 2.3–3.7% less diesel than the point-forecast controller and realises 76–86% of the achievable gain. Neural-network and gradient-boosting forecasters did not reduce fuel use further; planning around the forecast error did. Including grid pre-charging, 78–83% of the fuel energy of the no-storage case is saved. The benefit of choosing the operating point depends on the assumed low-load fuel curve.

Keywords: diesel generator; part-load efficiency; battery energy storage; model predictive control; forecast uncertainty; primary energy; residential microgrid; Iraq

Check for updates

Academic Editor: Luisa F. Cabeza

Received: 26 October 2026
Revised: 13 November 2026
Accepted: 20 November 2026
Published: 23 November 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Residential electricity supply in Iraq remains constrained by a persistent gap between generation and demand. In many districts of Baghdad and other cities, the public grid is delivered on a rotating timetable of a few hours on followed by a few hours off, adding up to roughly twelve hours of supply per day in summer, and the remainder is covered by neighbourhood or household diesel generators [1–3]. Fuel for such generators is a policy variable in its own right; for example, the price of diesel sold to private generator operators was halved for a three-month period in 2026 [4]. Backup generation during outages is a recurrent topic wherever supply is unreliable, from household battery-versus-generator comparisons [5,6] to the sizing of PV–battery–diesel systems for homes, telecom sites and rural areas [7–9]. The scale of this private generation, and the fuel it burns, makes the

efficiency with which small diesel sets convert fuel into electricity a question of national relevance as well as household economics.

Small diesel generators are poorly matched to residential demand. A household load is dominated by a low base demand with short peaks, so a set sized for the peak spends most of its running time at a small fraction of rated output. In this region the brake thermal efficiency (BTE) of a diesel engine falls and the specific fuel consumption (SFC) rises steeply, because friction and pumping losses and the fuel needed to keep the engine turning over do not scale with the delivered power [10,11]. Manufacturer data illustrate the effect: for an air-cooled three-cylinder DEUTZ generator engine, the SFC rises from 218 g/kWh at full load to 360 g/kWh at 25% load, an increase of about 65% [12]. Below 25% load the penalty is larger still, but it is rarely documented, and the quantity that drives the result, the no-load fuel flow, is usually not published.

Battery energy storage and photovoltaic (PV) generation are the obvious complements to such a generator. In hybrid-system design tools and optimisation studies, storage is credited mainly with avoiding capital and fuel cost, with the generator represented by an intercept-and-slope fuel curve [13–15]. Energy-management research for microgrids, from rule-based dispatch to model predictive control (MPC) and learning-based schedulers, is extensive [16–20], and scenario-based and chance-constrained predictive controllers have been proposed to handle forecast uncertainty [21–23]. Degradation-aware battery models and techno-economic assessments are available to refine the economics [24–27]. In most of this work the objective is monetary cost, and the saving is reported in currency or litres. What is less often examined is the thermodynamic mechanism: when storage lets a generator run fewer hours but at higher load, the fuel saved may exceed the energy that the battery actually delivers, because the avoided fuel was burned at poor efficiency.

Table 1 places the present study among the closest related work. Design and sizing studies of PV–battery–diesel systems for homes, telecommunication sites and off-grid communities [3,7–9,15,27] choose the component sizes for given demand profiles, and dispatch studies [19] operate a given plant; Ref. [20] addresses frequency support rather than energy dispatch. The question of how a diesel hybrid should be dispatched is older. Barley et al. [28] pointed out that the incentive for cycle charging, i.e., running the diesel at full power and storing the surplus, is its higher efficiency at full load, and found that simple dispatch rules with a frugal use of stored energy come close to an ideal predictive strategy. Aziz et al. [29] found for a PV/diesel/battery system that a combined dispatch rule reduced generator running time from 3675 to 2171 h per year, with 1.7% less fuel than load following and 24% less than cycle charging.

Predictive controllers that treat forecast uncertainty explicitly have been developed mainly for grid-connected microgrids. Alarcón et al. [21] combine an economic MPC that uses nominal forecasts with a scenario-based stochastic MPC that compensates forecast errors without assuming their distribution; Abdelghany et al. [22] integrate load forecasting with a scenario-based stochastic MPC for a PV–wind microgrid with hydrogen and battery storage that trades with the grid; and Babić et al. [23] build non-parametric chance constraints from kernel-density estimates of the forecast residuals of a PV–battery microgrid and report a cost up to about 1.5% below that of a nominal MPC. In these studies the objective is the monetary operating cost of a grid-connected system, and the fuel use of a diesel generator that depends on its operating point is not part of the decision. The authors' previous study [30] used the household data of the present work to train a deep Q-network that schedules the battery for cost and emission reduction against a rule-based baseline; the generator was not modelled as a fuel converter and forecast uncertainty was not treated explicitly.

Three questions therefore remain open for a household whose backup is a small diesel generator and whose outages follow a known timetable while its load is uncertain. First, how much of the fuel saved by a battery comes from displacing generator energy, and how much from moving the remaining generator operation to a more efficient operating point? Second, how much of the latter can a controller achieve without knowing the future load, and is a better forecasting model or an uncertainty-aware decision the more useful step? Third, does the saving survive a primary-energy balance that charges the grid energy used for pre-charging? The contribution of this study beyond [21–23,30] is not a new stochastic-control theory. It is the combination of an explicit part-load fuel model inside a scenario-based dynamic programme with generator top-up of forecast shortfalls, a set of reference strategies and bounds that separates the sources of the saving, and a thermodynamic and primary-energy accounting of the result.

Table 1. Position of this study among related work on storage and diesel dispatch.

Study	System	Decision Method	Forecast Uncertainty	Objective/Main Indicators
[28]	Wind–PV–diesel–battery, remote village	Dispatch rules compared with an ideal predictive strategy	Perfect-knowledge benchmark only	Fuel use, battery life, life-cycle cost
[29]	PV–diesel–battery	Sizing with load-following, cycle-charging and combined dispatch rules	Rules, no forecast	Cost, fuel use, generator hours
[21]	Microgrid with measured consumption	Economic MPC with a scenario-based stochastic MPC layer	Scenarios, no distribution assumed	Operating cost
[22]	PV–wind, hydrogen and battery storage, grid-connected	Scenario-based stochastic MPC with load forecasting	Scenarios of renewable output and load	Operating cost including storage degradation
[23]	PV–battery, grid-connected	Chance-constrained MPC	Non-parametric (kernel-density) chance constraints	Grid energy cost, constraint satisfaction
[30]	Same household, PV–battery, grid outages	Deep Q-network	Learned policy, no explicit forecast	Cost and emissions against a rule-based baseline
This work	PV–battery–10 kW diesel generator, scheduled outages	Scenario-based dynamic programme with generator top-up; point-forecast MPC; perfect-foresight bound	Three scenarios from out-of-fold error quantiles by lead time and time of day	Diesel fuel with part-load curve, grid energy, battery throughput; efficiency, primary energy, CO_2, cost

This paper treats the diesel generator as a thermal energy converter and assesses storage-enabled displacement on a fuel-energy and primary-energy basis. The contributions are as follows.

- A part-load fuel model for a small, ageing diesel set in Iraqi service, constructed from manufacturer SFC data with explicit, documented assumptions for the quantities that are not published (alternator efficiency, engine-age penalty, no-load fuel).
- A supply-timetable model for a Baghdad-type residential feeder, including random disturbance of the daily schedule, used to generate the generator duty for a full year of measured household load and PV data.
- A comparison of storage with grid pre-charge, a point-forecast receding-horizon controller, a scenario-based controller and a perfect-foresight bound, together with no-storage and rule-based references, which separates the fuel saving into the part due to

storage and grid pre-charge, the part due to choosing the generator operating point and the part lost to forecast error.

- A forecast-aware dispatch method that forms load scenarios from the out-of-fold quantiles of a ridge-corrected forecast and solves a scenario-based dynamic programme (DP) every hour with generator top-up of forecast shortfalls, and a test of whether neural and tree-based forecasters (multilayer perceptron (MLP), long short-term memory (LSTM) network, gradient boosting) improve the fuel result, evaluated both by forecast accuracy and by the fuel use of the controller that uses the forecast.
- A primary-energy balance that charges the energy drawn from the grid for pre-charging, together with CO_2 and abatement-cost metrics, and a sensitivity analysis on the uncertain generator parameters, battery size, fuel price and grid factor.

The same household data set was previously used to train a deep reinforcement-learning battery scheduler for cost and emission reduction [30]; here it serves as the load and PV input, while the question, the generator model, the dispatch method and the analysis are new.

The remainder of the paper is organised as follows. Section 2 describes the system, data, models, strategies, forecasters, evaluation indices and metrics. Section 3 presents the results in the order of the three questions above, Section 4 discusses them in relation to published work, together with their limitations, and Section 5 concludes.

2. Materials and Methods

Figure 1 summarises the study process and Figure 2 the energy flows of the system considered. The section follows the order in which the elements are used: the system and data (Sections 2.1–2.3), the generator and battery models (Sections 2.4 and 2.5), the dispatch strategies (Section 2.6), the forecasters and the scenario-based controller (Section 2.7), the indices used to evaluate forecasts and controllers (Section 2.8), the thermodynamic and economic metrics (Section 2.9) and the sensitivity analysis (Section 2.10).

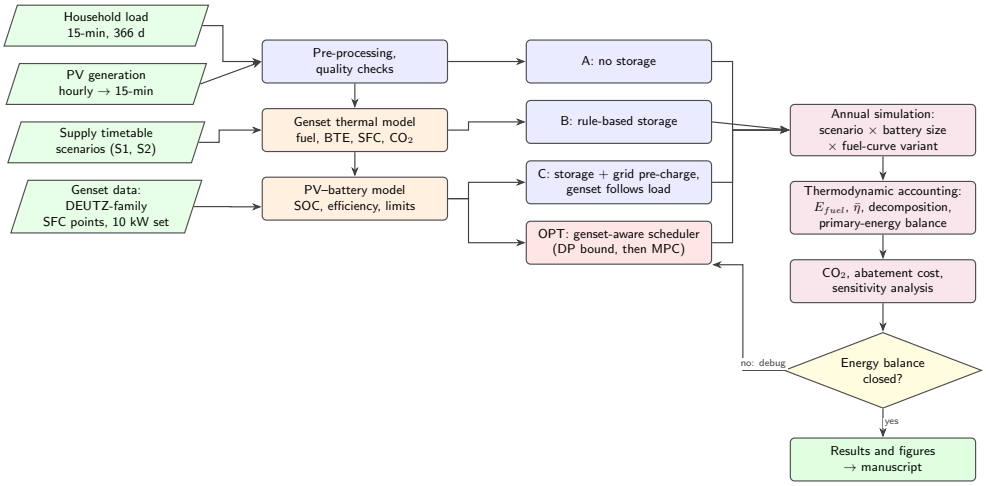


Figure 1. Overall study process: inputs, models, strategies and analysis.

2.1. System and Study Scope

A single grid-connected household in Iraq is considered, with a rooftop PV array, a hybrid inverter and a battery, and a privately owned diesel generator that supplies the household when the public grid is off (Figure 2). The household is used as a representative case; other households in the same region differ mainly in load magnitude, not in the shape of the profile, and a multi-household extension is left for future work. The generator is an

air-cooled two-cylinder set of the XHX-DA10 type, rated 10 kW electrical output with a DEUTZ F2L912-family engine of 14 kW listed power. Because units in service are older than the catalogue unit, the fuel curve is built from manufacturer data with an ageing penalty, as described in Section 2.4.

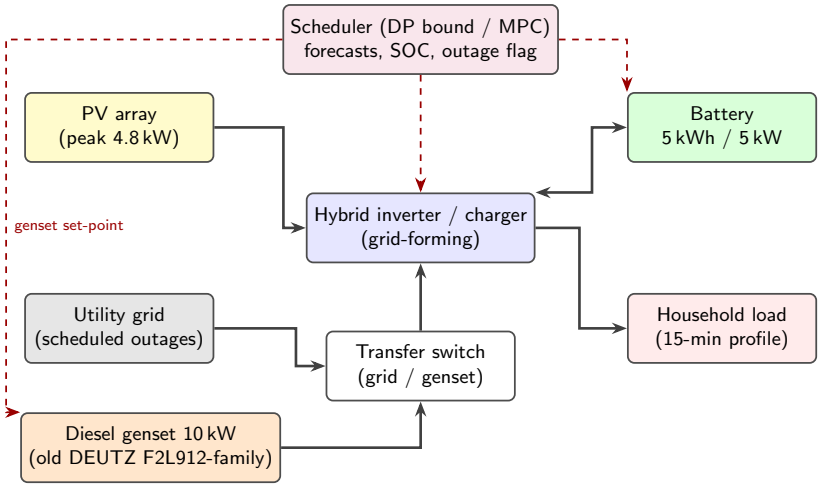


Figure 2. Energy-flow layout of the household system: PV, hybrid inverter, battery, grid, diesel generator with transfer switch, load and scheduler.

2.2. Data

The household load was recorded at 15 min resolution over one full year. The PV series was measured at hourly resolution at a real PV plant and represents the power the array can produce; the power actually used depends on the load, the battery state and the grid or generator state, and any surplus that cannot be used is curtailed in the model. The PV series contains 8784 hourly values. Each hourly value was taken as the power at the start of its hour and converted to the 15 min grid of the load by linear interpolation between consecutive hourly values; negative values were set to zero. The two series were aligned by their common start so that 15 min step k of the PV series and step k of the load series refer to the same time; the record is assumed to start at 00:00 on 1 January. The record contains 35,129 time steps of 15 min (365.9 days). The array has a peak output of about 4.8 kW. Ambient temperature is not used as a model input; the region experiences extreme summer temperatures, and derating of the generator and PV at high temperature is noted as a limitation and not modelled.

The public-grid timetable was not measured at the household. It is generated from a scenario model (Section 2.3). Data are in the repository folder `data/`(file `pvandpower.xlsx`).

2.3. Supply Timetable Scenarios

Two timetables are used, summarised in Table 2. Scenario S1 is a deterministic schedule built to reproduce the grid unavailability of the earlier study of this household [30], which reports 35.5% unavailability from May to August; the schedule used here gives 35.4%. In S1 the grid is off for 2 h in every 6 h on days 121–181 (1 May–30 June) and for 2 h in every 5 h on days 182–243 (1 July–31 August), and no interruption is applied in the other months, so the generator is required for 11.9% of the year. Scenario S2, the main scenario here, is a Baghdad-type rotation: in June to August the grid is on for three hours and off for three hours, starting at 06:00, giving 12 h of supply per day, with random disturbance. On 40% of days the whole pattern is shifted by ± 1 h, and on a further 40% of days one block is shortened by 1 h, so daily totals fall mostly between 10 and 13 h. This is consistent with reports of two-hour and three-hour rotations in Baghdad districts and longer cuts

elsewhere [1]; the pattern is a model and not a record of any particular feeder. In May and September a 2 h off, 4 h on pattern (8 h/day) is assumed; in October to April no scheduled interruptions are applied, which understates generator use and is a stated limitation. Over the year, the generator is required for about 17.7% of hours in S2. The stochastic scenario S2 is simulated with several random seeds (five for the strategy comparison, three for the controller comparison) and the mean is reported.

Table 2. Supply timetable scenarios. Percentages and the May/September pattern are assumptions.

Scenario	June–August	May, September	October–April
S1 (after [30])	2 h off per 6 h (May–June), 2 h off per 5 h (July–August)	included above	no interruptions
S2 (main)	3 h on/3 h off from 06:00, random disturbance	2 h off/4 h on	no interruptions

2.4. Diesel Generator as a Thermal Energy Converter

For electrical output P_{el} the shaft power is

$$P_{\mathrm{sh}} = \frac{P_{\mathrm{el}}}{\eta_{\mathrm{alt}}}, \tag{1}$$

with alternator efficiency η_{alt}. The fuel flow $\dot{F}$ (L/h) is a quadratic function of shaft power,

$$\dot{F}(P_{\mathrm{sh}}) = F_0 + c_1 P_{\mathrm{sh}} + c_2 P_{\mathrm{sh}}^2, \tag{2}$$

where F_0 is the no-load (idle) fuel flow. The fuel power, shaft efficiency, electrical efficiency and specific fuel consumption follow as

$$P_{\mathrm{fuel}} = \frac{\dot{F}\,\mathrm{LHV}}{3.6}, \tag{3}$$

$$\eta_{\mathrm{sh}} = \frac{P_{\mathrm{sh}}}{P_{\mathrm{fuel}}}, \qquad \eta_{\mathrm{el}} = \frac{P_{\mathrm{el}}}{P_{\mathrm{fuel}}} = \eta_{\mathrm{alt}}\,\eta_{\mathrm{sh}}, \tag{4}$$

$$\mathrm{SFC} = \frac{1000\,\rho\,\dot{F}}{P_{\mathrm{sh}}}, \tag{5}$$

with the lower heating value of diesel LHV = 35.8 MJ/L and fuel density $\rho = 0.835$ kg/L, and CO_2 emission $\dot{m}_{CO_2} = \varepsilon_{CO_2}\dot{F}$ with $\varepsilon_{CO_2} = 2.68$ kg/L. Over a period with generator operating steps k of duration Δt, the period-average electrical efficiency is

$$\bar{\eta}_{\mathrm{el}} = \frac{\sum_k P_{\mathrm{el},k}\,\Delta t}{\sum_k P_{\mathrm{fuel},k}\,\Delta t}. \tag{6}$$

This is the quantity that is lowered by prolonged low-load running.

Construction of the Fuel Curve

The published SFC points for the DEUTZ F3L912 genset engine [12] (Table 3) are applied as a shape to the 14 kW engine rating of the F2L912-family set, which is an assumption since the F2L912 table could not be retrieved. They are multiplied by an engine-age factor a (1.0, 1.15 or 1.30; default 1.15) to represent a worn engine, converted to fuel flow, and a quadratic through the four points and the idle fuel F_0 is fitted by least squares. The idle fuel is taken as a fraction of the fuel flow at rated output (0.20, 0.33 or 0.40; default 0.33), in line with the intercept-and-slope formulation used by HOMER [13], whose default intercept

(0.08145 L/h per kW rated) and slope (0.246 L/kWh) correspond to an idle fraction of about 0.25 [31]. For a 10 kW set the default curve of this study (idle 0.33, age 1.15) gives 1.5, 2.0, 2.7 and 3.4 L/h at 2.5, 5, 7.5 and 10 kW output, within about 5% of the HOMER default line at the three higher points; the low-age, low-idle limit (0.20, 1.0) reproduces the generic modern 20 kW consumption table scaled to 10 kW [32] to within 5% at the upper loads. The default therefore lies midway between a modern set and a worn one. Below 25% load the curve is therefore an extrapolation, and the idle fraction is the main uncertain input. It is varied in the sensitivity analysis and is to be replaced by a measured value (Section 2.10). Alternator efficiency is assumed to be 0.90.

Table 3. Specific fuel consumption points used to shape the fuel curve (DEUTZ F3L912 generator-set data [12]); the last column is after the default age factor of 1.15 (assumed).

Engine Load (%)	SFC, Datasheet (g/kWh)	SFC, Aged ($a = 1.15$) (g/kWh)
25	360	414
50	238	274
75	219	252
100	218	251

2.5. PV–Battery Model

The battery has nominal energy E_b and power limit $P_b^{max} = \min(5, E_b)$ kW (in kW, i.e., a C-rate of at most 1). The state of charge (SOC) s evolves as

$$s_{t+1} = s_t + \frac{\Delta t}{E_b}\left(\eta_c P_{c,t} - \frac{P_{d,t}}{\eta_d}\right), \qquad s_{min} \le s_t \le s_{max}, \tag{7}$$

with $P_c, P_d \ge 0$ the charge and discharge power at the DC bus, $\eta_c = \eta_d = 0.90$ and $s_{min} = 0.10$, $s_{max} = 0.90$. The main case is a 5 kWh, 5 kW battery, with 2.5 kWh/2.5 kW as a small-storage case and 10 kWh as a sensitivity case. At every step the power balance

$$P_{L,t} = P_{PV,t} + P_{g,t} + P_{grid,t} + P_{d,t} - P_{c,t} - P_{curt,t} \tag{8}$$

holds, where P_g is the generator electrical output and $P_{grid} = 0$ whenever the grid is off. Generator output above the load is allowed and goes to charging the battery.

2.6. Dispatch Strategies

Six strategies are compared: no storage (A), rule-based storage (B), storage with grid pre-charge (C), a point-forecast model predictive controller (MPC1), a scenario-based model predictive controller (SMPC) and a perfect-foresight optimal schedule (OPT). They differ in what the battery may be charged from, in how the generator output is chosen and in what the strategy knows about the future (Table 4). Three of them are reference strategies that follow fixed rules, two are controllers that could be deployed in a household and one is a bound that cannot be deployed. Comparing them in pairs isolates the sources of the fuel saving: B to C gives the effect of grid pre-charge, C to OPT the effect of choosing the generator operating point and C, MPC1 and SMPC against OPT the part of that effect that survives imperfect forecasts.

Table 4. Dispatch strategies. All strategies know the outage timetable of the current day; they differ in the information about the future load and PV.

Strategy	Battery Charged from	Generator Output	Knowledge of Future Load and PV	Role
A	– (no battery)	Follows the residual load	None	Reference without storage
B	PV surplus	Follows the load not covered by PV and battery	None	Rule-based battery, baseline of [30]
C	PV surplus; grid before each outage	As B	None	Grid pre-charge
MPC1	PV, grid, generator	Planned by the DP every hour; top-up of shortfalls	Point forecast	Deployable controller
SMPC	PV, grid, generator	Planned by the DP over three load scenarios every hour; top-up of shortfalls	Point forecast and error quantiles	Deployable controller (proposed)
OPT	PV, grid, generator	Planned by the DP over the whole year	Perfect	Upper bound

2.6.1. Reference Strategies

In strategy A there is no battery: PV serves the load directly, PV surplus is curtailed and during an outage the generator supplies the residual load. In strategy B, the rule-based battery used as the baseline in [30], PV surplus charges the battery and the battery discharges to cover any deficit; during an outage the generator supplies what PV and battery cannot. Strategy C adds grid pre-charge to B: whenever the grid is on before a scheduled outage, the battery is charged from the grid, so every outage starts with a full battery. In A, B and C the generator follows the net load, so it often runs at low output.

2.6.2. Perfect-Foresight Schedule (OPT)

In OPT the generator output is a decision variable. The generator may run above the load and charge the battery so that it runs fewer hours at a higher, more efficient output, or it may stay off. With full knowledge of load, PV and timetable over a horizon T, the schedule solves

$$\min_{\{P_{g,t},P_{b,t}\}} \sum_{t=1}^{T} \left[\dot{F}(P_{sh,t})\Delta t + \lambda\,|P_{b,t}|\,\Delta t + \lambda_g\,P_{grid,t}\,\Delta t \right] \tag{9}$$

subject to Equations (7) and (8), the SOC and power limits, $P_{g,t} \in \{0\} \cup [P_g^{min}, P_g^{max}]$ with $P_g^{max} = 10$ kW, and $P_{grid,t} = 0$ in outage steps. The first term is the diesel volume. The second term, with $\lambda = 10^{-4}$ L per kWh of battery throughput, discourages needless cycling. The third term expresses grid energy in litres of diesel at the price ratio of the two energy carriers, $\lambda_g = 0.035$ L/kWh (10 IQD/kWh at 1500 IQD/$, against 0.19 $/L of diesel; IQD: Iraqi dinar), so that grid energy is used when it is cheaper than generator fuel. The problem is solved by backward dynamic programming on a discretised SOC grid of 33 levels, a set of generator outputs (0 and 1.5 to 10 kW in 18 steps) and 11 battery power levels, with terminal value zero (Figure 3):

$$V_t(s) = \min_a \left[c(s,a) + V_{t+1}(s') \right], \tag{10}$$

where $c(s, a)$ is the bracketed stage cost of Equation (9) for action a in state s, and s' follows from Equation (7). A forward pass from the initial SOC then applies the stored optimal actions. Because the schedule uses perfect foresight, its fuel use is a lower bound that no deployable controller can beat.

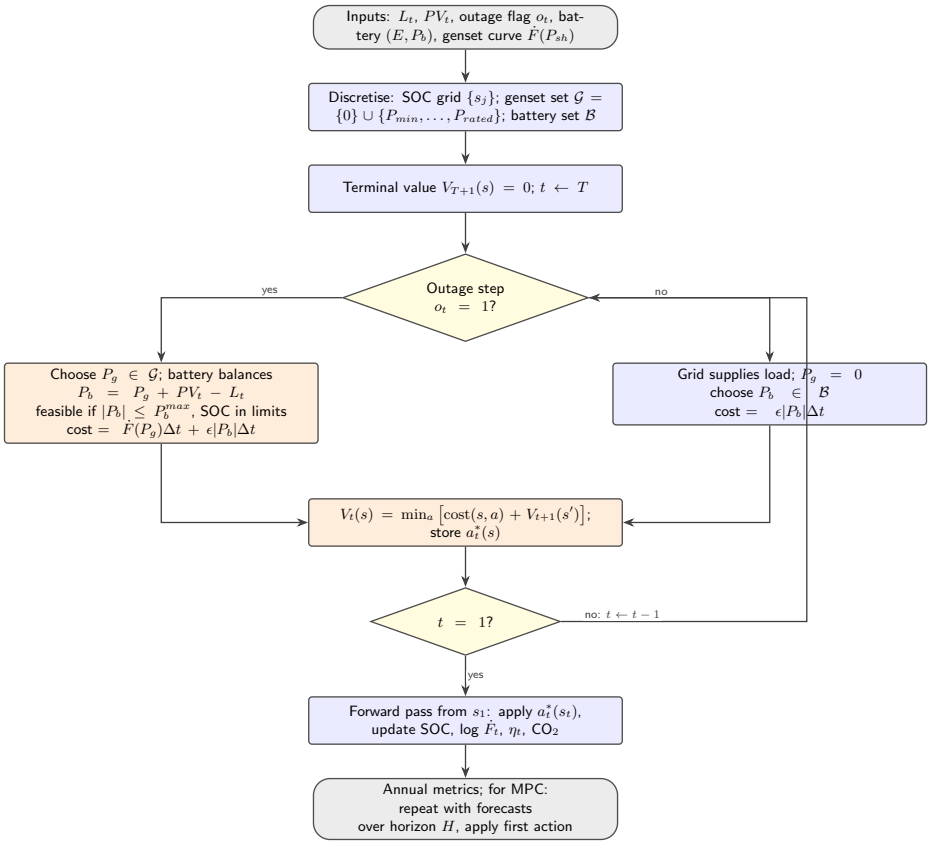


Figure 3. Dynamic programme used for the perfect-foresight schedule and, over a 24 h horizon, inside the controllers: backward recursion over outage and non-outage steps followed by a forward pass.

2.6.3. Deployable Controllers (MPC1 and SMPC)

The two deployable controllers use the same dynamic programme over a 24 h horizon, but with forecasts instead of the actual load and PV, and they repeat the optimisation every hour (receding horizon). MPC1 uses a single point forecast. SMPC, the controller proposed here, evaluates every candidate action over three load scenarios built from the forecast and the quantiles of its past errors. In both controllers the generator tops up any shortfall that arises when the actual load exceeds the forecast so that no load is left unserved. The forecasts, the scenarios and the controller are described in Section 2.7.

2.7. Forecasting, Scenarios and Scenario-Based Dispatch

This section describes how the forecasts used by MPC1 and SMPC are produced (Sections 2.7.1 and 2.7.2), how they were trained and validated (Section 2.7.3), how the load scenarios are formed (Section 2.7.4) and how the controller uses them (Section 2.7.5). Table 5 lists all settings.

Table 5. Settings of the forecasters and of the dispatch methods.

Component	Setting
Time step, horizon	15 min; 96 steps (24 h); re-solved every 4 steps (1 h)
Seasonal-naive forecast	Mean of the same time of day over the previous 7 days
Ridge correction	Inputs: deviation at t and at $t-4$; one coefficient set per lead time $(1,\dots,96)$; penalty $10^{-3}N$; eight contiguous folds
MLP	Inputs (296): 96 past loads, 96 past deviations, 96 seasonal-naive values of the horizon, 8 calendar values; two hidden layers of 128 units (ReLU); output: 96 corrections; 66,912 parameters
LSTM	Encoder LSTM (64 units) over past load and deviation (96 steps); decoder LSTM (64 units) with input seasonal-naive value, 4 calendar values and encoder state at each lead time; linear output; 52,033 parameters
MLP and LSTM training	Adam, learning rate 2×10^{-3}, batch 128, mean-squared-error loss; validation set: every sixth week of the training folds; early stopping with patience 8 epochs, at most 60 epochs; best-validation weights kept; seed 0
Gradient boosting	Histogram-based, pinball loss for $q=0.1, 0.5, 0.9$; 150 iterations, learning rate 0.1, at most 31 leaves; 9 features (Section 2.7.2); trained on every second origin
Cross-validation	Ridge, error quantiles and MATLAB controllers: eight contiguous blocks; machine-learning forecasters and Python controllers: five blocked folds by week; 8591 hourly origins
Error quantiles	10th, 50th, 90th percentiles; per lead time and six 4 h time-of-day bins; out-of-fold; bins with fewer than 20 samples use all bins
Scenarios	Three, weights 0.25/0.5/0.25; one action common to all scenarios
Dynamic programme	33 SOC levels (0.10–0.90); generator 0 and 1.5–10 kW in 18 steps; 11 battery power levels; terminal value $0.35\,(s_{\max}-s)E_{\mathrm{b}}$ L (controllers) or 0 (OPT); $\lambda=10^{-4}$ L/kWh; $\lambda_{\mathrm{g}}=0.035$ L/kWh; penalty on unserved energy 10^{3} L/kWh

2.7.1. Seasonal-Naive and Ridge-Corrected Forecast

The controllers in the main (MATLAB R2025a) simulation use a forecast of the load (and, separately, of PV) x for lead times $k=1,\dots,96$ steps (24 h). It starts from the seasonal-naive value $\bar{x}$, the mean of the same time of day over the previous seven days, and adds a ridge-regression correction from the two latest deviations from that value,

$$\hat{x}_{t+k|t}=\bar{x}_{t+k}+b_{0,k}+b_{1,k}\,(x_t-\bar{x}_t)+b_{2,k}\,(x_{t-4}-\bar{x}_{t-4}), \tag{11}$$

with one set of coefficients $(b_{0,k},b_{1,k},b_{2,k})$ for every lead time. The coefficients minimise the squared error plus a penalty of $10^{-3}N$ on their squared magnitude, where N is the number of training samples. The correction captures the persistence of the recent deviation, which decays with lead time. All forecasts are evaluated out of sample: the year is split into eight contiguous blocks, and the forecasts in one block use coefficients trained on the other seven.

2.7.2. Machine-Learning Forecasters

Three machine-learning forecasters, which are common alternatives to Equation (11), were tested in Python 3.14.6 64-bit (scikit-learn and PyTorch). All three predict the 96 future

loads from the same information as Equation (11) plus a longer history. Let $a_j = x_j - \bar{x}_j$ denote the deviation of the load from its seasonal-naive value, and $\mathbf{c}_j$ four calendar features of step j (sine and cosine of the hour of day and of the day of week).

Multilayer perceptron. The MLP is a feed-forward network that maps the input vector

$$\mathbf{z}_t = [x_{t-95}, \dots, x_t,\ a_{t-95}, \dots, a_t,\ \bar{x}_{t+1}, \dots, \bar{x}_{t+96},\ \mathbf{c}_{t+1},\ \mathbf{c}_{t+48}] \in \mathbb{R}^{296} \tag{12}$$

through two hidden layers to a correction of the seasonal-naive forecast,

$$\mathbf{h}_1 = \mathrm{ReLU}(\mathbf{W}_1\mathbf{z}_t + \mathbf{b}_1), \quad \mathbf{h}_2 = \mathrm{ReLU}(\mathbf{W}_2\mathbf{h}_1 + \mathbf{b}_2), \quad \hat{x}_{t+k|t} = \bar{x}_{t+k} + [\mathbf{W}_3\mathbf{h}_2 + \mathbf{b}_3]_k, \tag{13}$$

with 128 units per hidden layer, $\mathrm{ReLU}(u) = \max(u, 0)$ and 66,912 trainable parameters. The network therefore learns how the past day's load and deviations change the next day's profile, all lead times at once.

Long short-term memory network. The LSTM forecaster is an encoder–decoder. The encoder reads the past day, $\mathbf{u}_j = [x_j, a_j]$ for $j = t-95, \dots, t$, with the standard LSTM cell

$$\begin{aligned} \mathbf{i}_j &= \sigma(\mathbf{W}_i\mathbf{u}_j + \mathbf{U}_i\mathbf{h}_{j-1} + \mathbf{b}_i), & \mathbf{f}_j &= \sigma(\mathbf{W}_f\mathbf{u}_j + \mathbf{U}_f\mathbf{h}_{j-1} + \mathbf{b}_f), \\ \mathbf{o}_j &= \sigma(\mathbf{W}_o\mathbf{u}_j + \mathbf{U}_o\mathbf{h}_{j-1} + \mathbf{b}_o), & \mathbf{g}_j &= \tanh(\mathbf{W}_g\mathbf{u}_j + \mathbf{U}_g\mathbf{h}_{j-1} + \mathbf{b}_g), \\ \mathbf{c}'_j &= \mathbf{f}_j \odot \mathbf{c}'_{j-1} + \mathbf{i}_j \odot \mathbf{g}_j, & \mathbf{h}_j &= \mathbf{o}_j \odot \tanh(\mathbf{c}'_j), \end{aligned} \tag{14}$$

where $\mathbf{i}, \mathbf{f}, \mathbf{o}$ are the input, forget and output gates, $\mathbf{c}'$ the cell state, σ the logistic function and $\odot$ the element-wise product. The final encoder state $\mathbf{h}_t$ (64 units) summarises the past day. A second LSTM (64 units), the decoder, then runs over the 96 lead times with input $[\bar{x}_{t+k}, \mathbf{c}_{t+k}, \mathbf{h}_t]$, and a linear output layer gives the correction, $\hat{x}_{t+k|t} = \bar{x}_{t+k} + \mathbf{w}^\top\mathbf{d}_k + b$, where $\mathbf{d}_k$ is the decoder state. The model has 52,033 trainable parameters. Unlike the MLP, it processes the history in time order and can represent dependence on the sequence of recent events.

Gradient boosting. Histogram-based gradient boosting builds an additive model of regression trees, $F_M(\mathbf{v}) = F_0 + \nu \sum_{m=1}^{M} f_m(\mathbf{v})$, where each tree f_m is fitted to the negative gradient of the loss of the current model and ν is the learning rate. With the pinball loss $\rho_q(e) = \max\{q\,e, (q-1)\,e\}$ the model estimates the q-quantile of the target, so three models give $q = 0.1, 0.5, 0.9$ directly. The target is the deviation a_{t+k}, and each origin–lead pair is one sample with nine features: the lead time k, the hour and weekday of the target step, the seasonal-naive value $\bar{x}_{t+k}$, the latest deviation, the deviation one hour earlier, the mean deviation over the last two hours, the mean load over the last 24 h and the hour of the origin. The median model ($q = 0.5$) serves as a point forecast and the 10th and 90th percentile models as an alternative source of scenarios.

2.7.3. Training and Validation

Forecasts are issued every hour (every fourth step) from day 8 of the record, which gives 8591 forecast origins, each with 96 targets. For the machine-learning models, the origins are divided into five blocked folds of consecutive weeks; the model used for a fold is trained only on the other four, so every reported error and every forecast passed to a controller is out of sample. Within the training folds, every sixth week is held out as a validation set. The MLP and LSTM minimise the mean squared error over the 96 outputs with the Adam optimiser (learning rate 2×10^{-3}, batch size 128), with early stopping: training ends when the validation error has not improved for eight epochs (at most 60), and the weights of the epoch with the lowest validation error are kept. Inputs are in kW without further scaling, and the random seed is fixed (seed 0). Gradient boosting uses 150 iterations, a learning rate of 0.1 and at most 31 leaves per tree, and it is trained on every second origin to limit memory use.

The forecasters enter the dispatch study in two ways. In the main (MATLAB) simulation over the full year, MPC1 and SMPC use the ridge-corrected forecast of Equation (11). In a Python replication of the controllers over the outage season (May–September, 5 kWh battery, three timetable seeds), the point-forecast controller was run in turn with the seasonal-naive, ridge, MLP, LSTM and gradient-boosting (median) forecasts, and the scenario controller with scenarios from the ridge errors, the MLP errors and the gradient-boosting quantiles. At every hour the controller receives the out-of-fold forecast issued at that hour, so no forecast uses data from its own fold. This replication measures the effect of the forecaster on fuel use, which is the quantity of interest, rather than only its effect on forecast error.

2.7.4. Scenario Construction

The out-of-fold forecast errors $e = x - \hat{x}$ of the ridge forecast are grouped by lead time k and by the 4 h time-of-day bin b of the target step, and their 10th, 50th and 90th percentiles Q_{10}, Q_{50}, Q_{90} are computed from the other folds only. The three load scenarios are

$$L_{t+k}^{(q)} = \max\{\hat{L}_{t+k|t} + Q_q(k,b),\, 0\}, \qquad q \in \{10, 50, 90\}, \quad w = (0.25,\, 0.5,\, 0.25), \tag{15}$$

where the weights w approximate the probability mass that each quantile represents. Conditioning on the time-of-day bin lets the spread widen in the evening, when the load is most variable. The empirical quantiles need no assumption about the error distribution, which is skewed and heavy-tailed for a single household (Section 3.4).

2.7.5. Scenario-Based Receding-Horizon Dispatch

At every hour the dynamic programme of Section 2.6 is solved backward over the 24 h horizon with one action common to all scenarios. During an outage step the action is the planned generator output P_g; in scenario q the battery discharges first, $P_d^{(q)} = \min\{(L^{(q)} - P_{PV} - P_g)^+, P_b^{max}, (s - s_{min})E_b\eta/\Delta t\}$, and the generator tops up any remaining shortfall, $P^{top,(q)} = \min\{(L^{(q)} - P_{PV} - P_g - P_d^{(q)})^+, P_g^{max} - P_g\}$. The Bellman recursion is

$$V_t(s) = \min_a \sum_q w_q \Big[\dot{F}\big(P_{sh}(P_g + P^{top,(q)})\big)\Delta t + \lambda\,(P_c^{(q)} + P_d^{(q)})\Delta t + V_{t+1}\big(s'^{(q)}\big)\Big], \tag{16}$$

so a plan that leaves the battery nearly empty when the load may exceed the forecast is penalised by the fuel that the generator would burn at a low, inefficient output to cover the shortfall. On grid steps the stage cost is the throughput penalty plus the grid term of Equation (9). The terminal value $0.35\,(s_{max} - s)E_b$ litres charges the plan for energy that it leaves unused at the end of the horizon. With a single scenario (the ridge forecast) the same recursion gives MPC1. If the horizon contains no outage, the battery is charged at full power. The first four actions (one hour) are applied to the actual load and PV, and the procedure repeats (Figure 4). The scenario noise is treated as independent between steps, which is an approximation. The use of scenarios to represent forecast uncertainty follows the scenario-based predictive controllers of [21,22]; here it is combined with an explicit part-load fuel model of the generator and a top-up of forecast shortfalls.

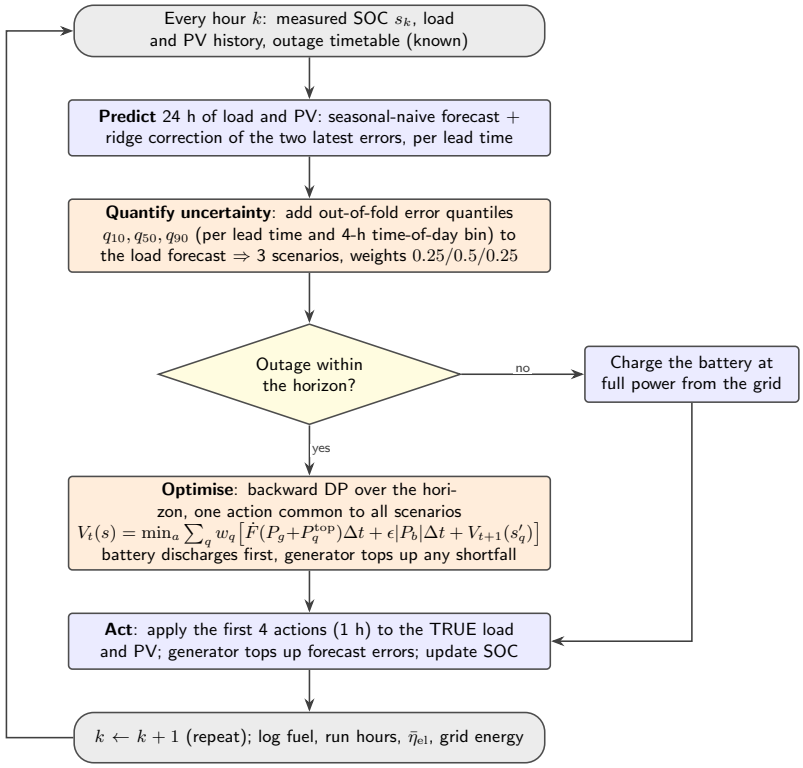


Figure 4. Forecast-aware dispatch: forecast, scenario construction from out-of-fold error quantiles, scenario-based dynamic programme (DP) and receding-horizon application.

2.7.6. Model and Controller Settings

Table 5 lists the settings of all forecasters and of the dynamic programme, and Table 6 in Section 2.9 the physical and economic parameters. The settings of the neural networks were fixed before the comparison and were not tuned on the test folds; early stopping on a separate validation set controls overfitting (Section 3.4).

Table 6. Main model parameters. "A" marks an assumption to be replaced by measurement.

Parameter	Value	Source
Generator electrical rating	10 kW	user data
Engine power listed	14 kW	user data
LHV of diesel	35.8 MJ/L	standard
Diesel density	0.835 kg/L	[12]
CO_2 factor	2.68 kg/L	standard
Alternator efficiency	0.90	A
Engine age factor	1.0/1.15/1.30	A
Idle fuel fraction	0.20/0.33/0.40	A
Battery energy/power	5 kWh/5 kW (2.5, 10 kWh)	design
Battery efficiency (each way)	0.90	A
SOC window	10–90%	A
Battery cost/life/rate	250 $/kWh, 10 y, 8%	A
Grid tariff, blocks per month	10/35/80/120 IQD/kWh	[33]
Exchange rate	1500 IQD/$	A
Time step	15 min	data

2.8. Statistical Indices for Forecast and Controller Evaluation

2.8.1. Forecast Accuracy Indices

For a forecast issued at origin t for lead time k, the error is $e_{t,k} = x_{t+k} - \hat{x}_{t+k|t}$. Over the set of N out-of-fold origins, the root-mean-square error at lead k, the mean absolute error over all leads and the skill score are

$$\mathrm{RMSE}_k = \sqrt{\frac{1}{N}\sum_t e_{t,k}^2}, \qquad \mathrm{MAE} = \frac{1}{96N}\sum_t \sum_{k=1}^{96} |e_{t,k}|, \qquad \mathrm{SS} = 1 - \frac{\mathrm{RMSE}}{\mathrm{RMSE}_{\mathrm{naive}}}, \tag{17}$$

where RMSE without subscript is taken over all leads and $\mathrm{RMSE}_{\mathrm{naive}}$ is that of the seasonal-naive forecast. RMSE and MAE are in kW. RMSE weights large errors quadratically, which matters here because a large under-forecast during an outage forces the generator to start at a low output; MAE gives the typical error and is less affected by occasional spikes. Their ratio indicates the shape of the error distribution: it is about 1.25 for normally distributed errors and larger for heavy-tailed errors. A positive skill score means that the model is more accurate than the seasonal-naive forecast. For the scenarios, the coverage of the 10–90% band, i.e., the fraction of observations that fall between $\hat{x} + Q_{10}$ and $\hat{x} + Q_{90}$ (prediction-interval coverage probability), is compared with its nominal value of 80%: a lower coverage means that the scenarios understate the risk of a high load, a higher one that they are conservative.

2.8.2. Controller Performance

The controllers are compared on the annual diesel volume F, the generator running hours, the period-average electrical efficiency of Equation (6) and the relative reduction of diesel use of strategy X against strategy Y, $R_{X/Y} = 100\,(F_Y - F_X)/F_Y$. Because grid pre-charge (C) and the perfect-foresight schedule (OPT) bracket what a controller can add by choosing the generator operating point, the share of that gain realised by a controller X is

$$S_X = 100\,\frac{F_{\mathrm{C}} - F_X}{F_{\mathrm{C}} - F_{\mathrm{OPT}}}, \tag{18}$$

which is 0% for a controller no better than grid pre-charge and 100% for one that matches perfect foresight. Results for the stochastic timetable are given as the mean over the timetable seeds, and the range over seeds is reported where it matters for the conclusion.

2.9. Thermodynamic and Economic Metrics

For each strategy and scenario the annual generator fuel volume $F = \sum_k \dot{F}_k \Delta t$, fuel energy $E_{\mathrm{fuel}} = F\,\mathrm{LHV}$, generator electricity E_{g}, running hours, $\bar{\eta}_{\mathrm{el}}$ from Equation (6) and CO_2 are reported. Savings relative to a reference strategy are

$$\Delta F = F_{\mathrm{ref}} - F_{\mathrm{str}}, \qquad \Delta E_{\mathrm{g}} = E_{\mathrm{g,ref}} - E_{\mathrm{g,str}}, \qquad \kappa = \frac{\Delta E_{\mathrm{fuel}}}{\Delta E_{\mathrm{g}}}, \tag{19}$$

where κ is the fuel energy saved per kWh of generator electricity displaced; its comparison with the average fuel energy per kWh in the reference indicates whether the saving exceeds the displaced energy. Because grid pre-charging transfers load to the grid, the net primary-energy balance includes the extra grid energy:

$$\Delta E_{\mathrm{p}} = \Delta E_{\mathrm{fuel}} - f_{\mathrm{g}}\left(E_{\mathrm{grid,str}} - E_{\mathrm{grid,ref}}\right), \tag{20}$$

where f_{g} is the primary-energy factor of grid electricity. No official value exists for Iraq, so a range is used: 2.0, 2.4 and 3.2. The middle value follows from the national generation mix (about 58% natural gas and 41% oil [34]) with typical gas- and oil-fired plant efficiencies and a 25% delivery loss, and 3.2 is a conservative case with higher losses; these are derived estimates, not published factors. ΔE_{p} is evaluated over a range of f_{g} and the break-even

value is reported. The corresponding CO_2 balance uses the grid emission factor in the same way; the sweep takes 0, 0.55 and 0.70 kg/kWh, bracketing the 0.549 kg/kWh reported for Iraq in 2023 (0.499 in 2022) [35].

The economics use the capital recovery factor $\mathrm{CRF} = r(1+r)^n/[(1+r)^n - 1]$ for discount rate r and life n years, giving an annualised battery cost $C_{\mathrm{ann}} = \mathrm{CRF}\, c_{\mathrm{b}} E_{\mathrm{b}}$ with specific cost c_{b}. The abatement cost is

$$\mathrm{AC} = \frac{C_{\mathrm{ann}} - p_{\mathrm{f}} \Delta F}{\Delta m_{\mathrm{CO_2}}}, \tag{21}$$

with fuel price p_{f} and $\Delta m_{\mathrm{CO_2}}$ the net CO_2 saving, and the numerator additionally includes the cost of the extra grid energy, $p_{\mathrm{e}}\,(E_{\mathrm{grid,str}} - E_{\mathrm{grid,ref}})$. The household pays the Iraqi increasing-block residential tariff [33]: 10 IQD/kWh for the first 1500 kWh of a billing period, 35 IQD/kWh for the next 1500 kWh, 80 IQD/kWh for the next 1000 kWh and 120 IQD/kWh above 4000 kWh (the blocks are assumed to apply per month). The same price applies to every kWh drawn from the grid, including battery charging, so grid pre-charging earns no time-of-use advantage; the bill is evaluated month by month from the simulated grid energy, and as long as monthly grid energy stays below 1500 kWh the marginal price is 10 IQD/kWh (≈0.0067 \$/kWh at 1500 IQD/\$, an assumed rate), which is also the price used in the scheduler's objective, Equation (9). The abatement cost is compared with estimates of the social cost of carbon [36] and read in the manner of marginal abatement cost curves [37]. Base assumptions are $c_{\mathrm{b}} = 250$ \$/kWh, $n = 10$ years, $r = 8\%$ and p_{f} between 0.19 and 0.50 \$/L, as the pump price in Iraq is about 288 IQD/L [38] and unsubsidised prices are higher; the battery cost is an assumption for an imported lithium iron phosphate unit installed by a local contractor: it lies above the global average turnkey price of large stationary systems of about 117 \$/kWh [39] and below the hardware cost alone (pack and inverter, about 470 \$/kWh) of a small United States residential system [40]. No survey of Iraqi household battery prices was found, so the cost is varied from 150 to 1000 \$/kWh in Section 3.11 and the break-even cost is reported in Table 7. Table 6 lists the main parameters.

Table 7. Installed battery cost (\$/kWh) at which the 5 kWh system breaks even (net annual benefit zero) and at which the abatement cost equals 35 and 44 \$/t, for the optimal schedule against no storage (10 years, 8%).

Diesel Price (\$/L)	Break-Even Cost	Cost for 35 \$/t	Cost for 44 \$/t
0.10	280	544	612
0.19	533	797	865
0.30	843	>1000	>1000
0.50	>1000	>1000	>1000

2.10. Design of the Sensitivity Analysis

The sensitivity analysis varies the idle fuel fraction (0.20–0.40), engine age factor (1.0–1.3), battery size (1.25–15 kWh), timetable scenario, fuel price, battery cost and grid primary-energy factor, one at a time and in the combinations where interaction is expected. A measurement of fuel flow at idle and at 25%, 50% and 100% load on the actual generator is planned and will replace the assumed curve.

3. Results

The results follow the three questions of Section 1. Sections 3.1–3.3 quantify the saving of the reference strategies and the perfect-foresight bound and separate its sources; Sections 3.4–3.6 show how much of the bound the deployable controllers achieve and what role the forecaster plays; Sections 3.7–3.11 give the primary-energy, CO_2 and economic balances, the trade-offs between objectives and the sensitivity to the uncertain inputs. All results are for the default fuel curve (idle fraction 0.33, age factor 1.15), the 5 kWh/5 kW battery as the main case and the S2 timetable unless stated. Energy balances close to numerical precision in every run (largest residual below 10^{-14} kW) and no load is left unserved.

3.1. Generator Efficiency over the Load Range

Figure 5 shows the assumed fuel curve. At the default setting the fuel flow is 1.5, 2.0 and 3.4 L/h at 2.5, 5 and 10 kW electrical output, corresponding to fuel-to-electricity efficiencies of 17%, 25% and 30%. The efficiency at the lowest part-load points therefore is about half of that at rated output, and the idle-fraction and age settings move the whole curve by about ±10 percentage points at rated output. Without storage the generator follows the household load and operates at an average efficiency of about 15% (S2: 3344 kWh from 2291 L); with rule-based storage the figure is 10–12%, because the battery removes the load peaks and leaves the generator with the low load only.

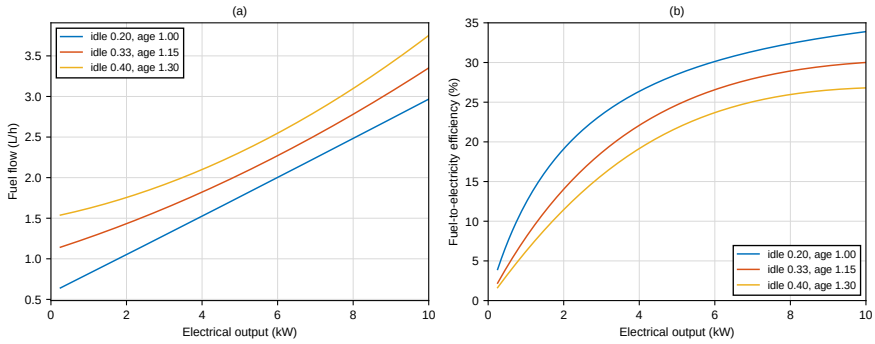


Figure 5. Assumed fuel flow (**a**) and fuel-to-electricity efficiency (**b**) of the 10 kW generator for three settings of the idle fuel fraction and engine age factor. The curve below 25% load is an extrapolation.

3.2. Annual Diesel Use by Strategy

Table 8 and Figure 6 give the annual results; the generator electricity and efficiency are given for every strategy, including C and the scenario controller. With no storage the generator burns 2291 L per year and runs through every outage hour (about 1550 h). For the 5 kWh battery, rule-based storage reduces diesel to 1203 L, grid pre-charge to 268 L and the optimal scheduler to 196 L; running time falls from 882 h (B) to 84 h (OPT). Diesel use falls steeply up to about 5 kWh and then flattens: the marginal saving per added kWh of battery drops from about 140 L (2.5 to 5 kWh) to about 23 L (7.5 to 10 kWh), and at 15 kWh the generator is hardly needed (3–6 kWh per year). In the deterministic timetable S1, the 5 kWh battery reduces diesel from 1512 L (A) and 777 L (B) to 71 L (OPT), so the same pattern holds under a different outage schedule.

Table 8. Annual diesel use, generator electricity and period-average electrical efficiency by strategy and battery size (S2 timetable). A, B, C and OPT: mean of 5 seeds. SMPC: mean of seeds 1–3, simulated for the two battery sizes of the controller comparison (Table 9); –: not simulated. A: no storage; B: rule-based storage; C: storage with grid pre-charge; SMPC: scenario controller; OPT: perfect-foresight schedule.

Battery (kWh)	A	B	C	SMPC	OPT
Diesel use (L/yr)					
1.25	2291	1397	1049	–	897
2.5	2291	1307	700	569	550
5	2291	1203	268	212	196
7.5	2291	1111	129	–	104
10	2291	1041	54	–	46
15	2291	998	3	–	3
Generator electricity (kWh/yr)					
1.25	3344	1614	1252	–	1339
2.5	3344	1532	920	1001	1001
5	3344	1383	447	529	515
7.5	3344	1234	226	–	262
10	3344	1121	98	–	114
15	3344	1046	4	–	6
Period-average electrical efficiency $\bar{\eta}_{el}$ (%)					
1.25	14.7	11.6	12.0	–	15.0
2.5	14.7	11.8	13.2	17.7	18.3
5	14.7	11.6	16.8	25.1	26.5
7.5	14.7	11.2	17.7	–	25.3
10	14.7	10.8	18.3	–	24.9
15	14.7	10.5	15.4	–	22.2

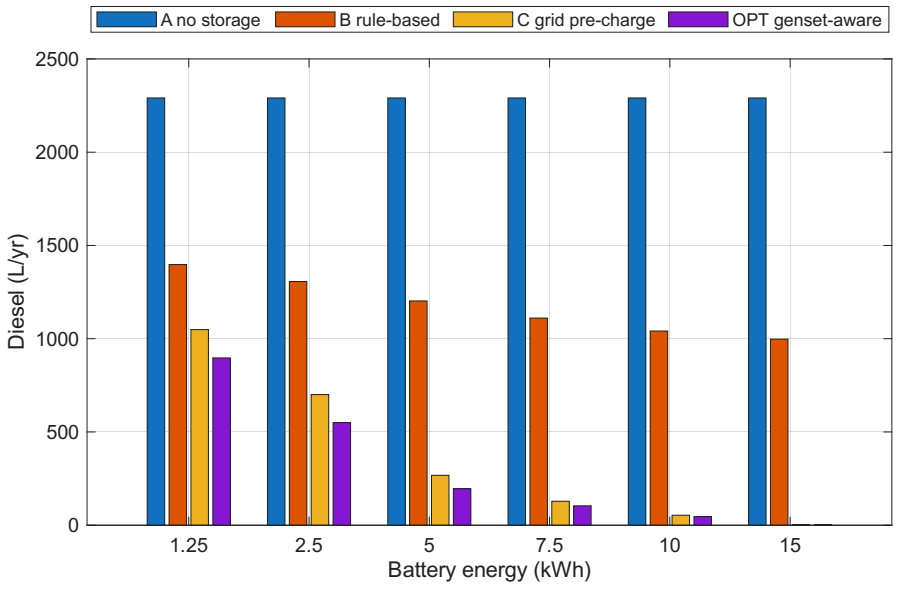


Figure 6. Annual diesel use by strategy and battery size (S2, mean of 5 seeds).

Table 9. Diesel use and generator operation of the baselines, the deployable controllers and the bound (S2, mean of seeds 1–3, battery power = $\min(5\ \text{kW}, E_b)$). A: no storage; B: battery with rule-based operation (baseline of [30]); C: grid pre-charge; MPC1: point-forecast controller; SMPC: scenario controller; OPT: perfect foresight. Share: fraction of the C-to-OPT diesel saving realised. $\bar{\eta}_{el}$: period-average electrical efficiency; run: generator hours per year.

Battery (kWh)	Quantity	A	B	C	MPC1	SMPC	OPT
2.5	Diesel (L/yr)	2292	1305	698	585	569	548
	Run (h/yr)	1555	953	493	375	356	340
	$\bar{\eta}_{el}$ (%)	14.7	11.8	13.2	17.0	17.7	18.3
	Share (%)	–	–	–	75.5	86.4	100
5	Diesel (L/yr)	2292	1200	266	219	212	195
	Run (h/yr)	1555	881	172	104	96	84
	$\bar{\eta}_{el}$ (%)	14.7	11.6	16.8	24.3	25.1	26.4
	Share (%)	–	–	–	66.6	76.4	100

3.3. Where the Saving Comes from

Table 10 separates the effects. Against rule-based storage, the optimal schedule cuts diesel by 84% and generator electricity by only 63% at 5 kWh; against no storage the figures are 91% and 85%. The fuel saving therefore exceeds the energy displaced, and the fuel energy saved per kWh of generator electricity displaced (7.4 kWh against 6.8 kWh average fuel use per kWh in the no-storage case) is higher than the average, because the energy removed is the least efficient part of the generator's operation. Most of the reduction comes from storage with grid pre-charge (A and B to C). Choosing the generator operating point adds a further 22% (2.5 kWh) and 27% (5 kWh) of diesel saving over C, but only 14% at 10 kWh and nothing at 15 kWh, where the battery covers every outage. The optimal schedule actually runs the generator for more kWh than C (515 against 447 kWh at 5 kWh; Table 8) but at much higher output and for far fewer hours, raising efficiency from 17% to 26% (Figure 7). The two effects should be read on different scales. Relative to the no-storage case, grid pre-charge alone already achieves 96.6% of the diesel reduction of the perfect-foresight schedule at 5 kWh ($(2291-268)/(2291-196)$), so the operating-point effect is 27% of the diesel that remains after pre-charge but about 3% of the original consumption, or 72 L per year.

Table 10. Reduction of diesel use and generator electricity (S2, mean of 5 seeds). The last column is the extra diesel saving of OPT over C, i.e., the effect of choosing the generator operating point.

Battery (kWh)	Diesel, OPT vs. A (%)	Diesel, OPT vs. B (%)	Generator kWh, OPT vs. A (%)	Generator kWh, OPT vs. B (%)	Diesel, OPT vs. C (%)
1.25	60.8	35.8	60.0	17.0	14.5
2.5	76.0	57.9	70.1	34.7	21.5
5	91.5	83.7	84.6	62.7	26.9
7.5	95.5	90.6	92.2	78.8	18.9
10	98.0	95.6	96.6	89.8	14.0
15	99.9	99.7	99.8	99.5	−0.5

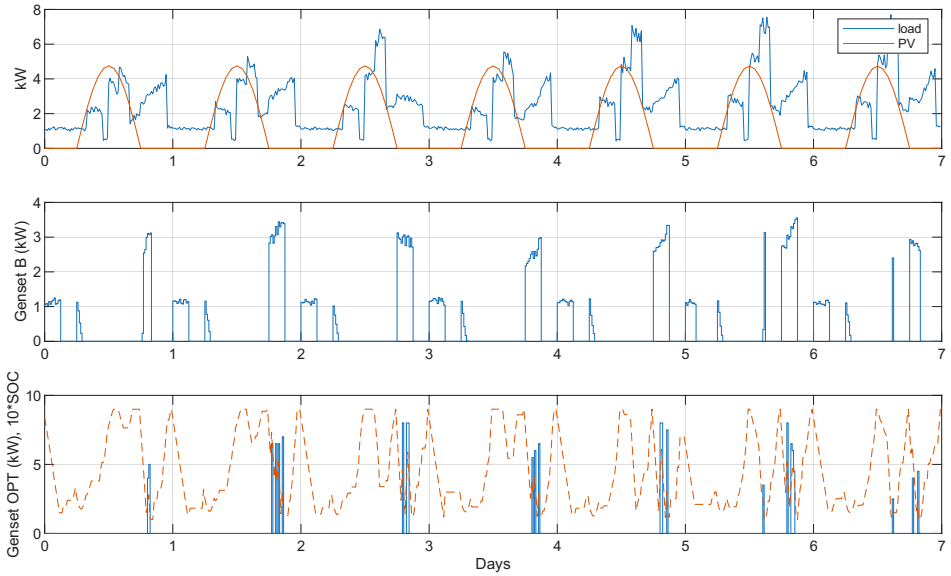


Figure 7. Example summer week: household load and PV (**top**), generator output under rule-based storage B (**middle**) and under the optimal schedule (**bottom**, solid) with battery SOC scaled by 10 (dashed).

3.4. Forecast Accuracy

3.4.1. Accuracy by Lead Time

Table 11 compares the load forecasts out of sample with the indices of Section 2.8. The seasonal-naive forecast has an RMSE of 0.356 kW at every lead time, against a mean load of 1.43 kW, since it does not use recent observations. The ridge correction and the MLP reduce the error to 0.24 kW one hour ahead, but beyond about 6 h every model is close to the naive value, which shows that the household load is dominated by a random component that no model removes. Over all leads the MLP is the most accurate model (skill score 11%), followed by the LSTM and the ridge correction (5%) and gradient boosting (1%). The MLP is 6–8% more accurate than the ridge correction from 3 to 12 h; the LSTM is better than the ridge correction only at 3 and 6 h and is the least accurate at 1 h. The ratio of RMSE to MAE is about 1.9 for every model, well above the value of 1.25 expected for normally distributed errors: the errors are heavy-tailed, with occasional large under-forecasts when appliances are switched on, which is why the scenarios are taken from empirical quantiles rather than from a normal distribution. The 10–90% band of the ridge-error quantiles covers 78.5% of the observations against a nominal 80%, while that of the gradient-boosting quantile models covers 71.7%, so the latter understate the spread. The PV forecast error is only 0.03 kW because the PV record is the measured production of a plant with little day-to-day variability.

Table 11. Out-of-fold load-forecast accuracy: RMSE (kW) by lead time, RMSE and MAE (kW) over all leads, and skill score SS against the seasonal-naive forecast (Equation (17)); blocked five-fold cross-validation, 8591 origins; negative forecasts set to zero; the MLP and LSTM use early stopping on a validation set. Mean load: 1.43 kW.

Forecast	1 h	3 h	6 h	12 h	24 h	RMSE	MAE	SS (%)
Seasonal-naive	0.356	0.356	0.356	0.356	0.356	0.358	0.180	0
Ridge correction (Equation (11))	0.240	0.324	0.347	0.357	0.320	0.341	0.177	4.7
MLP	0.241	0.298	0.324	0.330	0.323	0.318	0.174	11.1
LSTM	0.277	0.307	0.337	0.363	0.326	0.340	0.179	4.9
Gradient boosting (median)	0.271	0.320	0.352	0.372	0.361	0.355	0.179	0.8

3.4.2. Training Behaviour of the Neural Forecasters

Figure 8 shows the training, validation and test error of the two networks on one fold. The training error of the MLP falls steadily, while its validation error reaches its minimum after five epochs and then rises slowly, the usual sign of overfitting; early stopping keeps the weights of epoch 5. The validation error of the LSTM is noisy and reaches its minimum after four epochs. In both cases the test error follows the validation error, so the networks were stopped at their best generalisation rather than under-trained, and neither approaches the accuracy needed to remove the random part of the load.

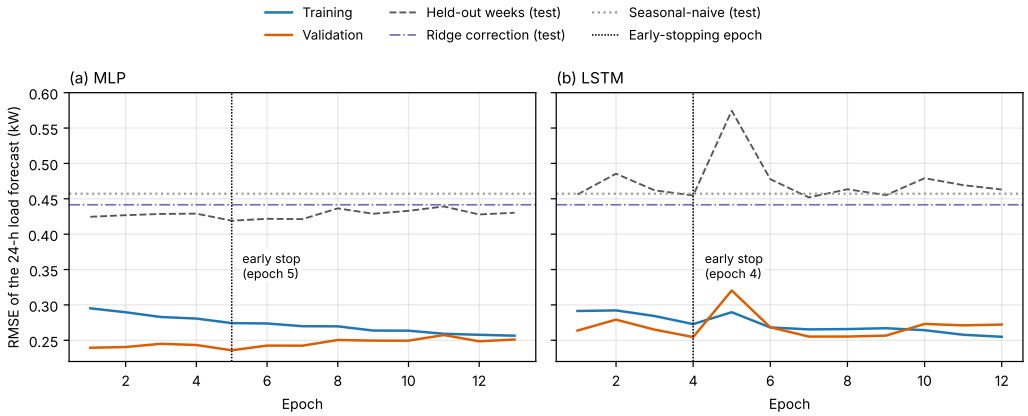


Figure 8. Training, validation and test error (RMSE, kW) per epoch of the MLP and the LSTM on one held-out fold (5903 training, 1008 validation and 1680 test samples), with the epoch of minimum validation error marked. The dashed lines are the ridge correction and the seasonal-naive forecast on the same test weeks.

3.5. Scenario-Based Dispatch

3.5.1. Comparison of the Controllers

Table 9 and Figure 9 give the diesel use of the deployable controllers against the bounds. At 5 kWh the scenario controller SMPC uses 212 L per year, 20.5% less than grid pre-charge (C) and 3.17% less than the point-forecast controller MPC1; it realises 76% of the C-to-OPT saving, against 67% for MPC1. At 2.5 kWh the corresponding figures are 569 L, 18.5%, 2.77%, 86% and 76%. SMPC is better than MPC1 for every seed and battery size, by 2.3–3.7% (range over the three seeds and two sizes), and it runs the generator for fewer hours (96 against 104 at 5 kWh) at a higher average efficiency (25.1% against 24.3%). Compared with the battery operated by simple rules (B), the scheduled battery cuts the generator running time from 881 to 96 h per year at 5 kWh (89%) and from 953 to 356 h at 2.5 kWh (63%), and it reduces the generator electricity by 62% and 35%. Against the earlier baselines, SMPC reduces diesel use by 90.8% (5 kWh) and 75.2% (2.5 kWh) relative to no storage and by 82.3% and 56.4% relative to rule-based storage. The remaining distance to the perfect-foresight bound is 8.7% of the diesel at 5 kWh and 3.7% at 2.5 kWh, which is the cost of the unpredictable part of the load. No unserved energy occurred and the energy balance closed to 10^{-15} kW. For reference, the forecast-light rule and the point-forecast controller of the first formulation, which did not allow generator top-up inside the optimisation, used 236 and 219–221 L at 5 kWh.

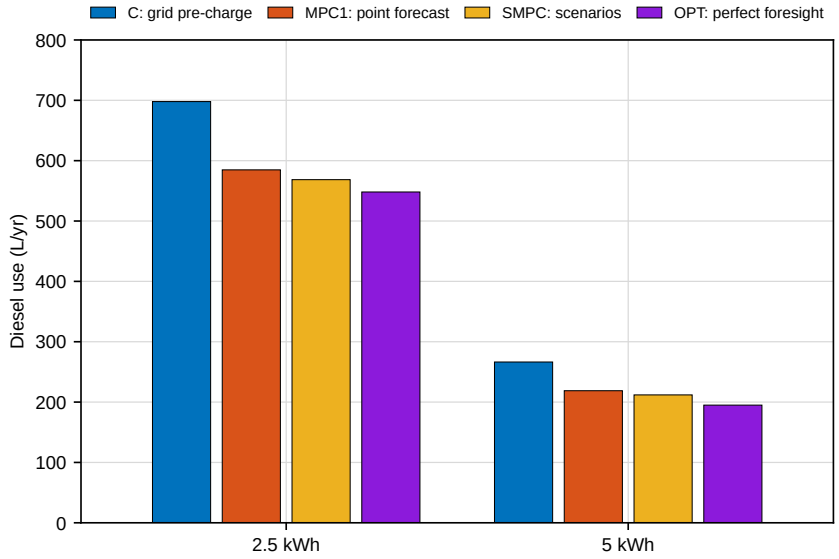


Figure 9. Annual diesel use of grid pre-charge (C), the point-forecast controller (MPC1), the scenario controller (SMPC) and the perfect-foresight bound (OPT) for two battery sizes (S2, full year, mean of seeds 1–3; error bars: minimum and maximum over the seeds).

3.5.2. Effect of the Forecasting Model on Fuel Use

In the Python replication of the controllers over the May–September season (5 kWh, three timetable seeds; Section 2.7.3), a point-forecast controller on the seasonal-naive forecast realised 68.9% of the C-to-OPT saving, on the ridge forecast 82.2%, on the gradient-boosting median 77.2%, on the early-stopped MLP forecast 79.2% and on the early-stopped LSTM forecast 76.4%. The MLP has the smallest error between 3 and 12 h, yet it realised less of the saving than the ridge forecast in two of the three seeds. Forming scenarios from the ridge forecast raised the share to 89.3% and from the MLP forecast to 86.8%, and using gradient-boosting quantiles for the spread gave 87.5%. The benefit comes from allowing for the forecast error in the decision, not from a more complex forecast model, and the simple empirical quantiles of the ridge correction were the best choice. The ranking by fuel use therefore differs from the ranking by RMSE in Table 11: an 11% skill score did not translate into a lower fuel use.

3.6. Time-Resolved Behaviour

Figure 10 shows three consecutive days of the closed-loop simulation (5 kWh, one timetable seed, 3 h on/3 h off supply). Grid pre-charge (C) keeps the battery full whenever the grid is on and empties it in the evening outage, so the generator starts at the end of the outage once the battery is exhausted. The scheduled controllers behave differently in three ways: they let the state of charge fall to 10–30% outside the outages and refill it from PV and from the grid in stages, instead of holding it full; and they concentrate the generator output into short bursts at 6–9 kW, where its efficiency is highest, instead of long periods at low load. The scenario controller differs from the point-forecast controller mainly in its reserve: it holds a higher state of charge (about 0.3 against 0.1 at hours 3–12 of the first day) in anticipation of a load above the forecast, which is what reduces the number of generator starts. The battery power of the receding-horizon controllers changes sign often within an outage period because the plan is recomputed every hour; this chatter is small in energy terms but would matter for battery ageing and is the reason for the throughput penalty.

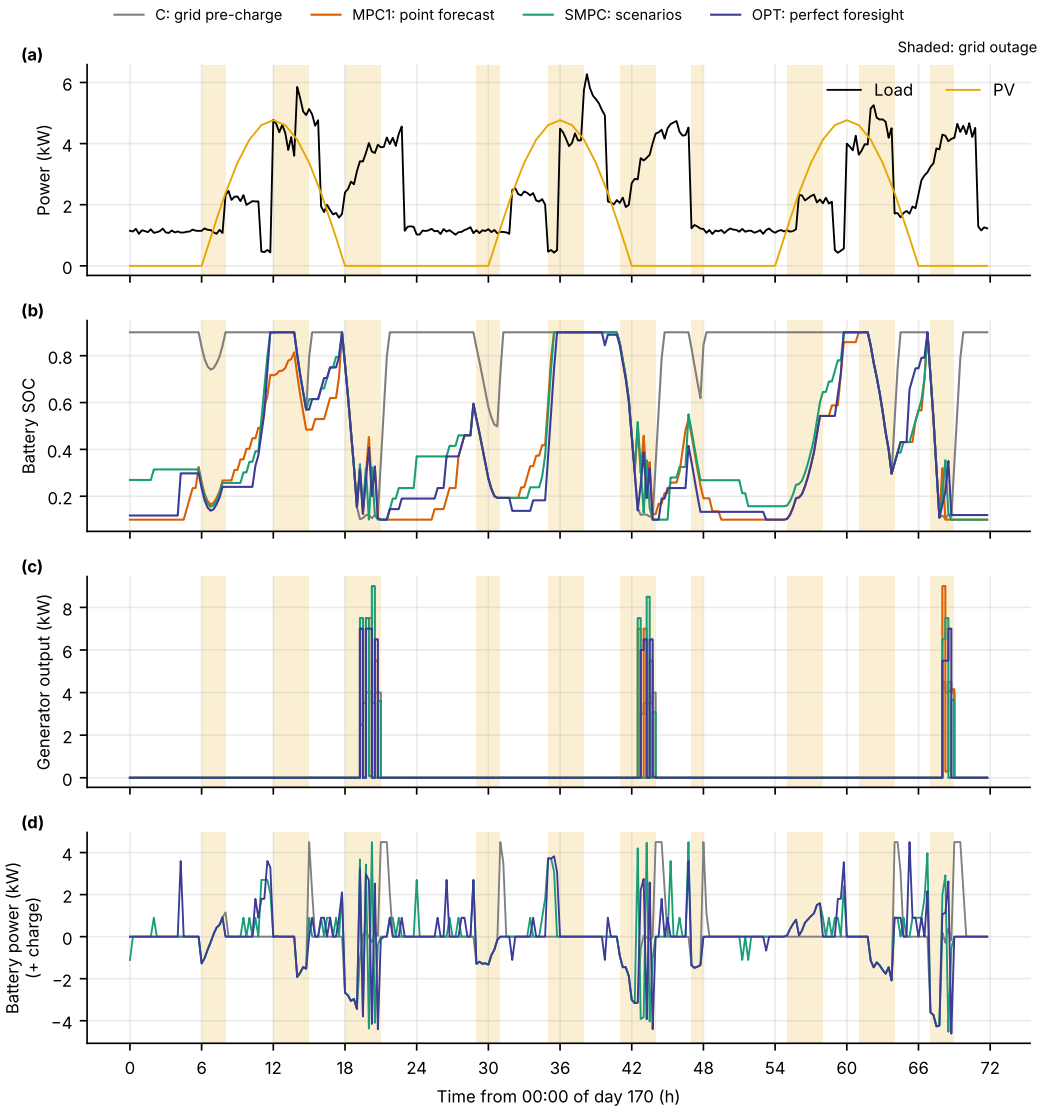


Figure 10. Three days of closed-loop operation in the outage season (5 kWh, seed 11, Python replication): (**a**) household load and PV power with the grid outages shaded; (**b**) battery state of charge; (**c**) generator electrical output; (**d**) battery power. C: grid pre-charge; MPC1: point-forecast controller; SMPC: scenario controller; OPT: perfect foresight.

Over the outage season (days 121–273) the cumulative diesel use of the controllers separates mainly after day 150, when the supply changes to 3 h on/3 h off (Figure 11). The excess over the perfect-foresight bound grows steadily at a nearly constant rate for each controller, which indicates that the gap is a systematic cost of not knowing the load rather than a few bad days: by the end of the season it is 69 L for C, 11 L for MPC1 and 7 L for SMPC.

3.7. Primary Energy and CO_2

Grid pre-charge shifts load from the generator to the grid, so the net primary-energy balance of Equation (20) is evaluated against the grid primary-energy factor f_g. For batteries up to 5 kWh, the extra grid energy is only 126–216 kWh per year, against about 1400–2100 L of diesel saved, and the fuel energy saved is 80 to 155 times the extra grid energy (break-even f_g of 80–155, against $f_g \leq 3$ in the sweep). For 7.5 kWh and above, the grid energy actually falls, because stored PV energy that the no-storage case curtails is used, so the net saving

(Figure 12) exceeds the fuel saving. The CO_2 saving versus no storage is 4.7 t (2.5 kWh) and 5.6 t (5 kWh) per household per year for a zero grid emission factor and changes by less than 2% for emission factors up to 0.7 kg/kWh at 5 kWh; at the reported Iraqi emission factor of 0.549 kg/kWh [35] the 5 kWh saving is about 5.5 t.

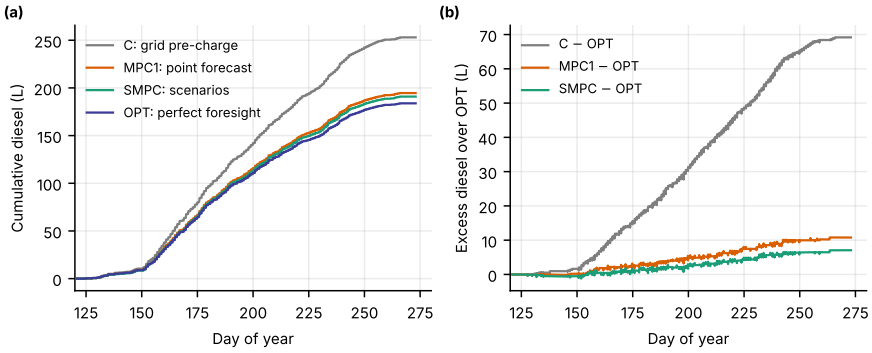


Figure 11. (**a**) Cumulative diesel use and (**b**) excess over the perfect-foresight schedule during the outage season (5 kWh, seed 11, Python replication).

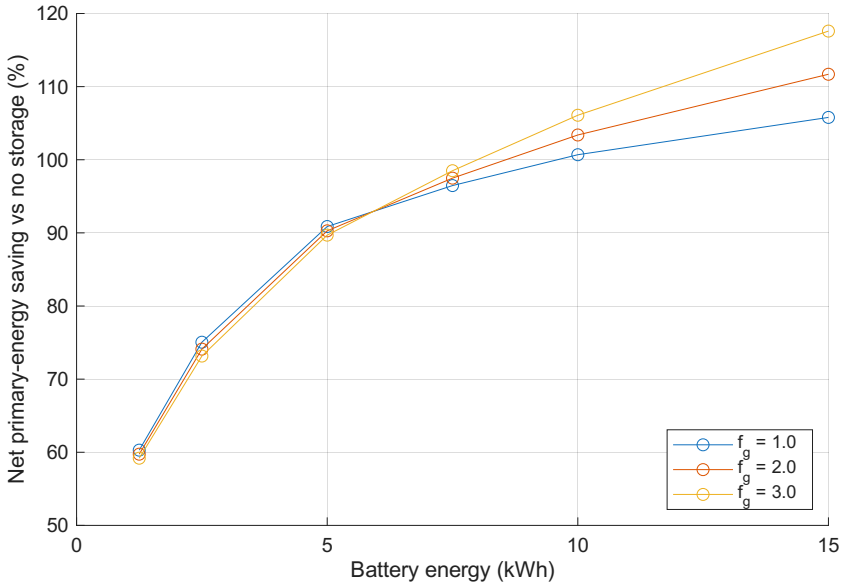


Figure 12. Net primary-energy saving of the optimal schedule relative to the fuel energy of the no-storage case, for grid primary-energy factors f_g of 2.0, 2.4 and 3.2, and the fuel-energy saving alone (dashed). Values above 100% occur when grid purchases fall as well (S2, mean of 5 seeds).

3.8. Economics

Table 12 gives the economics at 250 $/kWh. The extra electricity bill under the Iraqi block tariff is under 2 $ per year at 5 kWh, because monthly grid energy stays below 710 kWh and the 10 IQD/kWh block applies throughout. The net annual benefit is positive for batteries up to 10 kWh at the pump price of 0.19 $/L and up to 15 kWh at 0.50 $/L, with a maximum at 2.5–5 kWh; at 15 kWh and 0.19 $/L it is negative. Under these assumptions the abatement cost is negative (about −38 $/t at 5 kWh), i.e., the system pays for itself. An abatement cost of 35–44 $/t at 5 kWh corresponds to an all-in installed cost of about 800–865 $/kWh (battery, inverter, installation, and replacement), which is plausible for a household installation but has to be argued, not assumed.

3.9. Benefit Realised by the Scenario Controller

The primary-energy, CO_2 and economic results above were computed for the perfect-foresight schedule and are upper bounds. For the scenario controller the same balances were evaluated directly in the simulation against the no-storage case on the same timetable (S2, mean of seeds 1–3; Table 13). At 5 kWh SMPC saves 2080 L of diesel per year (20.7 MWh of fuel energy) and draws 881 kWh more from the grid; the extra grid energy is larger than for OPT (135 kWh) because SMPC does not time the use of stored PV energy as well, but the fuel energy saved is still 23.5 times the extra grid energy, so the net primary-energy balance stays positive for any grid primary-energy factor below 23.5 and equals 78–83% of the fuel energy of the no-storage case for factors of 2.0–3.2. The net CO_2 saving is 5.09 t per household per year at the reported Iraqi grid emission factor of 0.549 kg/kWh [35] (5.57 t at zero and 4.96 t at 0.70 kg/kWh); OPT gives 5.55 t at 0.549 kg/kWh. At 250 $/kWh, 8% and 10 years, the net annual benefit is 203 $ at a diesel price of 0.19 $/L and 848 $ at 0.50 $/L, which is 96% and 98% of the perfect-foresight benefit, and the abatement cost is −40 $/t. The monetary difference between SMPC and the point-forecast controller is small, about 1 $ per year at 0.19 $/L and 3 $ at 0.50 $/L (7 L of diesel), so the value of the scenario formulation lies in the fuel and run-hour reduction and in the proximity to the bound, not in the household bill.

Table 12. CO_2 and economics of the OPT schedule versus no storage (S2, mean of 5 seeds). Net annual benefit = diesel saved × fuel price − extra grid bill − annualised battery cost (250 $/kWh, 10 years, 8%). The CO_2 saving is for zero grid emission factor; the abatement cost is negative where the system pays for itself.

Battery (kWh)	CO_2 Saved (t/yr)	Extra Grid Energy (kWh/yr)	Net Benefit at 0.19 $/L ($/yr)	Net Benefit at 0.50 $/L ($/yr)	Abatement Cost at 0.19 $/L ($/t)
1.25	3.7	126	217	650	−58
2.5	4.7	216	236	776	−51
5	5.6	134	211	860	−38
7.5	5.9	−232	138	815	−23
10	6.0	−615	58	754	−10
15	6.1	−1345	−115	594	19

Table 13. Primary energy, CO_2 and economics of the deployable controllers against no storage (S2, mean of seeds 1–3, simulated; 250 $/kWh, 8%, 10 years). PE: net primary-energy saving as a share of the fuel energy of the no-storage case, for grid primary-energy factors of 2.0, 2.4 and 3.2. CO_2 at 0.549 kg/kWh. Benefit and abatement cost at the fuel prices stated.

Battery (kWh)	Controller	Diesel Saved (%)	Extra Grid (kWh)	PE Saving (%)	CO_2 (t/yr)
2.5	MPC1	74.5	622	69.0/67.9/65.8	4.23
	SMPC	75.2	627	69.7/68.6/66.4	4.27
	OPT	76.1	217	74.2/73.8/73.0	4.55
5	MPC1	90.5	878	82.8/81.2/78.1	5.07
	SMPC	90.8	881	83.0/81.5/78.4	5.09
	OPT	91.5	135	90.3/90.1/89.6	5.55
Battery (kWh)	**Controller**	**Benefit at 0.19 $/L**	**Benefit at 0.50 $/L**	**Abatement ($/t)**	**Break-Even f_g**
2.5	MPC1	227	756	−54	27.3
	SMPC	230	764	−54	27.4
	OPT	237	777	−52	79.8
5	MPC1	202	844	−40	23.5
	SMPC	203	848	−40	23.5
	OPT	211	861	−38	154.2

3.10. Multi-Objective Trade-Offs

The dispatch problem has three competing objectives: diesel use F, grid energy E_g and battery throughput Q. The cost function used so far is the weighted sum $J = F + \lambda_g E_g + \lambda_w Q$. Varying one weight at a time, with the other at its reference value, traces the trade-off (5 kWh, one timetable seed, full year; Table 14). The weighted-sum method finds only points on the convex part of the Pareto front, so the table shows a lower bound on the available trade-offs, not the whole front.

Table 14. Multi-objective sweep for the perfect-foresight schedule (OPT) and the scenario controller (SMPC), 5 kWh, one timetable seed, full year. E_g is the annual grid import, Q the battery throughput. Reference: $\lambda_g = 1 \times 0.035$ L/kWh, $\lambda_w = 10^{-4}$ L/kWh.

Varied	Value	OPT			SMPC		
		F (L)	E_g (kWh)	Q (kWh)	F (L)	E_g (kWh)	Q (kWh)
λ_g	0×	192.6	7008	3283	199.6	7001	3300
	0.25×	192.9	6013	5054	205.2	6767	3372
	1×	192.3	6008	5054	206.3	6756	3368
	4×	192.2	5996	5054	209.2	6733	3348
	16×	1048	3837	6578	1048	4629	4869
	256×	1199	3694	6837	1221	4468	5180
λ_w (L/kWh)	0	192.6	6008	5431	207.5	6761	3644
	10^{-3}	192.7	6008	5044	206.9	6754	3358
	10^{-2}	193.1	6018	4976	207.7	6755	3304
	3×10^{-2}	192.6	6767	3281	207.7	6764	3284
	10^{-1}	204.5	6765	3130	220.8	6757	3105

3.10.1. Weight on Grid Energy

Multiplying λ_g by 0.25–4 changes diesel use by less than 0.4% (192.2–192.9 L for OPT), so the reference weight lies on a flat part of the front; setting it to zero raises grid import by about 1000 kWh (17%) with no fuel saving in these runs. Weights of 16 times the reference or more make grid energy more expensive than diesel and the generator replaces the grid: diesel use rises to 1048 L and run time from 84 to 437 h while grid import falls by about 36%. There is no intermediate compromise because the generator's fuel cost per kWh at part load exceeds the grid weight by a large factor; the trade-off is a step, not a smooth curve, and the reference weight is the correct side of it.

3.10.2. Weight on Battery Throughput

The throughput weight shows a real knee: raising λ_w from 0 to 3×10^{-2} L/kWh cuts the annual battery throughput of the perfect-foresight schedule by 40% (5431 to 3281 kWh) at the same fuel use (192.6 L), because much of the cycling at zero weight has no fuel value; the price is 760 kWh (13%) more grid import, since part of the cycling was used to shift grid energy. Beyond that the battery is no longer used enough and diesel rises (204 L at 0.1 L/kWh). The reference value of 10^{-4} L/kWh is therefore very small; a value of about 3×10^{-2} L/kWh would reduce cycling by 40% at no fuel cost but with more grid import, and it is a better setting only if battery lifetime is valued above that energy; this should be decided with a measured cycle-ageing cost. The scenario controller follows the same pattern but at 4–8% higher diesel use; its throughput is lower than that of OPT at the reference setting (3300 against 5050 kWh) because it cannot exploit the information that causes OPT to cycle.

3.11. Sensitivity Analysis

3.11.1. Generator Fuel Curve

Table 15 and Figure 13 show that the main conclusion does not depend on the assumed fuel curve. Across all nine combinations of idle fraction (0.20–0.40) and age factor (1.0–1.3), the 5 kWh battery with the optimal schedule reduces diesel by 89.9–92.2% against no storage and by 79.6–85.4% against rule-based storage and generator electricity by 62.8–63.0% against rule-based storage. At 2.5 kWh the corresponding ranges are 74.5–76.8%, 52.7–60.2% and 34.6–35.2%. The two parameters act differently. The age factor scales all diesel figures (OPT: 170, 195 and 220 L at age 1.0, 1.15 and 1.3, idle 0.33) but leaves every relative saving unchanged. The idle fraction controls the value of choosing the operating point: the extra saving of OPT over C rises from 14.0% to 24.7% at 2.5 kWh and from 16.7% to 31.5% at 5 kWh as the idle fraction goes from 0.20 to 0.40, while the absolute diesel use of OPT at 5 kWh barely changes (194–196 L at age 1.15), because the generator then runs only about 84 h per year at high output. The efficiency of the generator under rule-based storage ranges from 9% to 17% and under OPT from 23% to 31% across the grid of curves, so the efficiency recovery is robust in direction but not in size.

Table 15. Sensitivity of the 5 kWh result to the assumed fuel curve and battery efficiency (S2, mean of seeds 1–3). The first nine rows vary idle fuel fraction and engine age factor at 90% battery efficiency; the default is idle 0.33, age 1.15.

Idle	Age	η_b	A (L)	B (L)	OPT (L)	OPT vs. A (%)	OPT vs. B (%)	OPT vs. C (%)
0.20	1.00	0.90	1693	837	171	89.9	79.6	16.7
0.20	1.15	0.90	1947	963	196	89.9	79.6	16.8
0.20	1.30	0.90	2201	1088	222	89.9	79.6	16.7
0.33	1.00	0.90	1993	1044	170	91.5	83.8	26.8
0.33	1.15	0.90	2292	1200	195	91.5	83.8	26.8
0.33	1.30	0.90	2591	1357	220	91.5	83.8	26.8
0.40	1.00	0.90	2154	1155	168	92.2	85.4	31.5
0.40	1.15	0.90	2478	1328	194	92.2	85.4	31.5
0.40	1.30	0.90	2801	1502	219	92.2	85.4	31.4
0.33	1.15	0.85	2292	1212	235	89.8	80.7	21.1
0.33	1.15	0.95	2292	1189	176	92.3	85.2	28.5

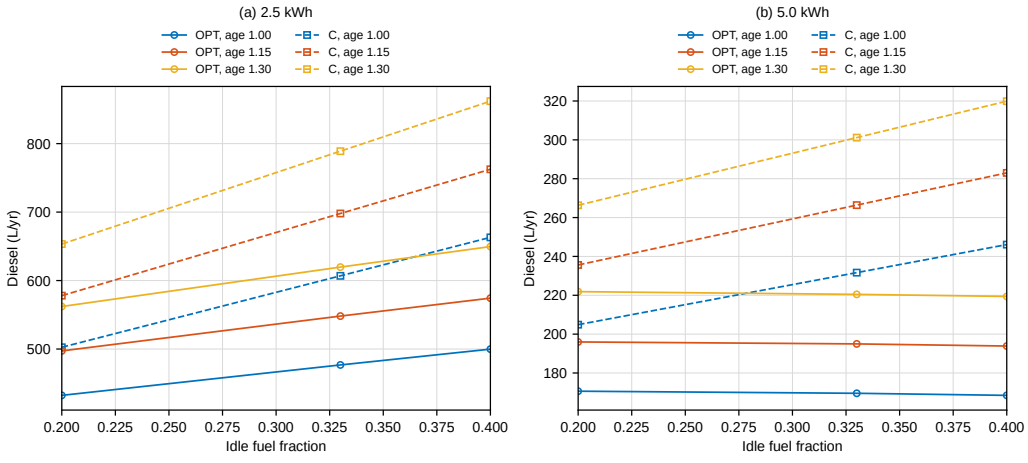


Figure 13. Annual diesel use of the optimal schedule (solid) and of storage with grid pre-charge C (dashed) against the idle fuel fraction, for three engine age factors.

3.11.2. Battery Efficiency

Lowering the round-trip efficiency from 0.81 (0.90 each way) to 0.72 (0.85 each way) raises OPT diesel at 5 kWh from 195 to 235 L and the extra grid energy from 135 to 273 kWh per year; raising it to 0.90 (0.95 each way) lowers diesel to 176 L and makes grid use fall slightly (−10 kWh). The reduction against no storage stays between 89.8% and 92.3%.

3.11.3. Fuel Price and Battery Cost

Figure 14 and Table 7 show the economics against fuel price and installed cost. At the pump price of 0.19 $/L, the 5 kWh system breaks even at an installed cost of about 530 $/kWh, and the abatement cost reaches 35–44 $/t at about 800–865 $/kWh; at 0.10 $/L those costs are 280 and 545–610 $/kWh, and at 0.50 $/L the system pays for itself even at 1000 $/kWh. The values apply to the optimal schedule; the implementable controllers use about 12% more diesel at 5 kWh, which changes the annual saving by less than 5 $ at 0.19 $/L.

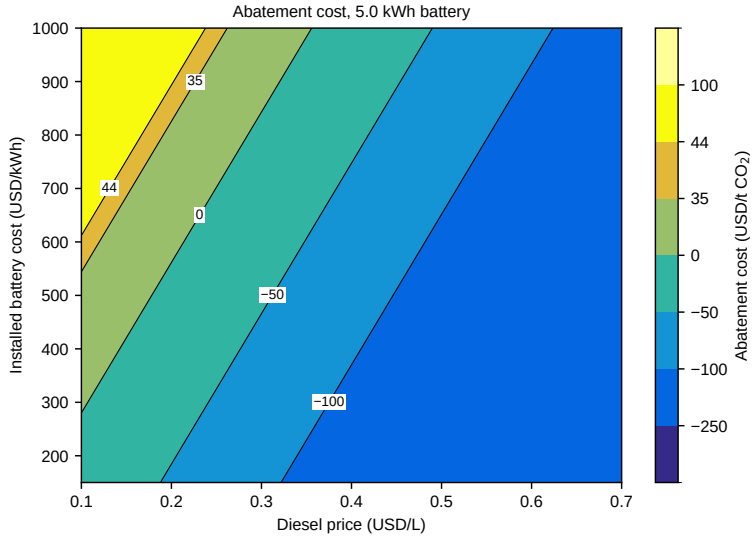


Figure 14. Abatement cost ($/t CO_2) of the 5 kWh system against diesel price and installed battery cost; negative values mean the system pays for itself.

3.11.4. Battery Size and Idle Fuel

Figure 15 combines the two assumptions that matter most, the battery capacity and the no-load fuel of the generator, for the perfect-foresight schedule (full year, seeds 1–2, S2 timetable). Battery size dominates: at the reference idle fraction of 0.33 the diesel saving against no storage rises from 45% at 1 kWh to 76%, 92%, 96% and 98% at 2.5, 5, 7.5 and 10 kWh, with diminishing returns above 5 kWh; against rule-based storage a 1 kWh battery saves only 11%, because it cannot cover an outage evening. The idle fraction changes the saving much less: at 5 kWh the saving against no storage is 89–93% for idle fractions of 0.15–0.50, and the diesel use of the scheduled battery barely moves (189–192 L), because the schedule avoids running the generator at low load. The saving against rule-based storage is more sensitive (78–87% at 5 kWh), since the rule-based battery still runs the generator at low load and therefore pays the idle fuel. The main conclusion, a saving of about 90% at 5 kWh, therefore does not depend on the least certain assumption of the generator model; the figures are upper bounds, which the scenario controller approaches to within the shares reported above.

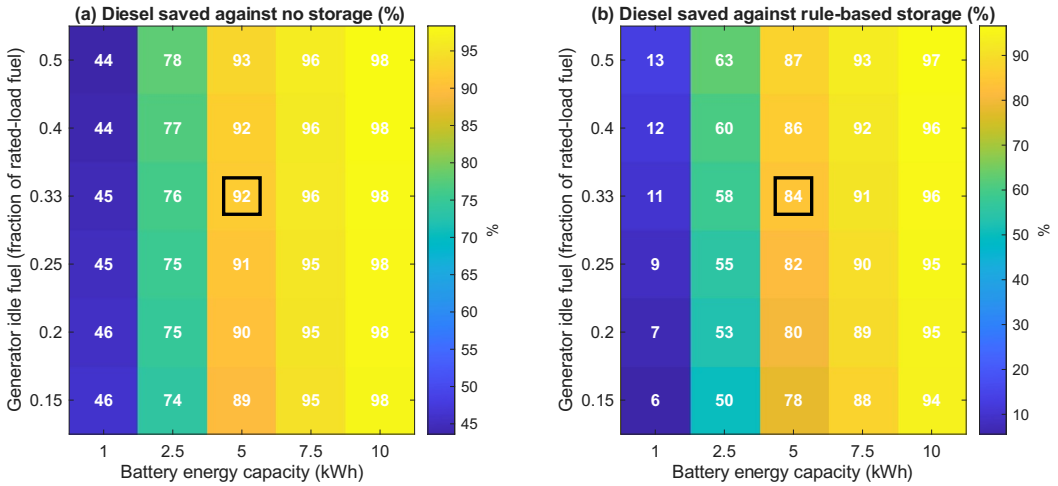


Figure 15. Diesel saving of the perfect-foresight schedule against (**a**) no storage and (**b**) rule-based storage, as a function of battery capacity and the generator idle-fuel fraction (full year, seeds 1–2). The box marks the reference case.

4. Discussion

This section interprets the results, compares them with published work and states where they apply. Sections 4.1–4.4 follow the three questions of Section 1, Section 4.5 discusses applicability beyond the case study and Section 4.6 lists the limitations.

4.1. Thermodynamic Mechanism of the Saving

The results support a thermodynamic reading of storage-enabled displacement. A diesel generator that follows a residential load operates where its electrical efficiency is 12–17%; a battery lets it either stop or run for short periods at 60–100% of rated output, where its efficiency is 26–30%. The fuel saved therefore exceeds the energy displaced: each kilowatt-hour of generator electricity removed by the optimal schedule carries 7.4 kWh of fuel energy, against an average of 6.8 kWh per kilowatt-hour in the no-storage case, because the kilowatt-hours that disappear are those produced at the lowest efficiency. This is the effect that Barley et al. [28] identified as the incentive for cycle charging in remote hybrid systems. The present results show that it is large in a household setting for a specific reason: the mean household load (1.43 kW) is only 14% of the generator rating, so without storage the generator spends almost all of its running time in the region where the fuel curve is dominated by the no-load fuel.

4.2. Sources of the Saving and Comparison with Dispatch Studies

Most of the diesel reduction comes from having stored energy available at the start of every outage, which is what grid pre-charge provides. Relative to no storage, grid pre-charge achieves 96.6% of the reduction of the perfect-foresight schedule at 5 kWh, which agrees with the finding of [28] that simple dispatch rules with a sensible use of stored energy come close to an ideal predictive strategy. Choosing the generator operating point adds a further 27% of the remaining diesel at 5 kWh and 22% at 2.5 kWh, falling to 14% at 10 kWh and to nothing at 15 kWh, where the battery covers every outage. This is a more conservative claim than attributing the whole saving to efficiency recovery: in absolute terms the operating-point effect is about 70 L per household per year at 5 kWh.

The direction of the operating-point effect agrees with Aziz et al. [29], where a combined dispatch rule reduced generator running time by 41% compared with load following but fuel use by only 1.7%. Here the perfect-foresight schedule reduces running time by 51%

(from 172 to 84 h at 5 kWh) and fuel use by 27% compared with grid pre-charge. The fuel effect is larger for two reasons: the household generator works at a much lower load factor, and the assumed idle fuel (one third of the full-load fuel flow) makes low-load running expensive. Table 15 confirms the second reason: the extra saving over grid pre-charge falls from 31.5% to 16.7% at 5 kWh when the idle fraction is reduced from 0.40 to 0.20. The operating-point effect is therefore the result that depends most on the measurement of the fuel curve, whereas the total saving against no storage (about 90% at 5 kWh) does not (Figure 15).

4.3. Value of Scenario-Based Control

The scenario controller needs only a seven-day load history, the known outage timetable and a small table of error quantiles, and it solves a 24 h dynamic programme in a fraction of a second per hour, so it can run on a household controller. It realises 76–86% of the gain that perfect foresight would deliver over grid pre-charge and uses 2.3–3.7% less diesel than a point-forecast controller built on the same forecast. The size of this gain is of the same order as that reported for other stochastic predictive controllers: Babić et al. [23] found a cost at most about 1.5% below that of a nominal MPC for a PV–battery microgrid, with the main benefit in constraint satisfaction rather than cost. The same pattern appears here. In money the difference between SMPC and MPC1 is small (about 7 L, or 1–3 $, per household per year at 5 kWh), and the benefit lies in robustness: SMPC keeps a larger reserve in the battery (Figure 10), runs the generator 8% fewer hours (96 against 104 h at 5 kWh) and is better than MPC1 for every seed. Unlike the grid-connected cases of [21–23], where a forecast error is absorbed by the grid at a known price, a forecast error during an outage here forces the generator to start at a low and inefficient output; this is the cost that the scenarios price, and it explains why an explicit fuel curve inside the optimisation is needed. Because the benefit saturates above about 5 kWh and the saving per added kilowatt-hour falls, a 2.5–5 kWh battery is the economically sensible range under the assumed prices.

4.4. Forecasting Versus Decision-Making

Neither the MLP, the LSTM nor gradient boosting improved the fuel result over a ridge-corrected seasonal forecast, although the MLP had the best forecast accuracy (skill score 11% against 5% for the ridge correction). This is not in conflict with studies that find LSTM networks more accurate than benchmark methods for individual households [41]: accuracy did improve here as well. What the results show is that the improvement in RMSE did not reach the decision. Two reasons are consistent with the evidence. First, the load of a single household at 15 min resolution is dominated by noise beyond a few hours (Table 11), and the review of low-voltage forecasting by Haben et al. [42] notes that household-level load is very volatile, that pointwise measures such as RMSE can be misleading at this level and that simple linear models are strong benchmarks. Second, the controller is sensitive to the tail of the error distribution during the outage evening, not to the average error over all hours, and the scenarios address that tail directly. The practical conclusion is that forecasters for this type of controller should be selected by the fuel use of the closed loop, not by RMSE alone, and that modelling the forecast error in the decision was the step that reduced fuel use.

4.5. Applicability Beyond the Case Study

Back-up generation of this kind is widespread: an estimated 20–30 million back-up generator sites in developing countries spend up to 50 billion $ a year on fuel and emit more than 100 Mt of CO_2 [43]. The mechanism studied here requires three conditions, and it should transfer wherever they hold: an outage pattern that is known or announced

in advance, a generator that is large compared with the load it serves and a battery that can supply a substantial part of the energy of an outage period. The relative savings depend mainly on the ratio of battery energy to the outage-evening demand (Figure 15: a 1 kWh battery saves only 11% against rule-based storage, a 5 kWh battery 84%) and on the idle fuel of the generator. Where the generator serves a flatter or larger load, as at the telecommunication sites of [8] or with neighbourhood generators that supply many households, the generator already runs closer to its rated output and the operating-point effect is expected to be smaller, although the displacement effect remains. The absolute figures in litres, hours and dollars apply to the case study only and should be recomputed with the local load, timetable, generator and prices.

4.6. Limitations

The limitations of the study are as follows.

(i) The fuel curve of the actual old generator is assumed from a manufacturer datasheet of a related engine, with an assumed age factor and idle fraction; a measurement is needed, and the operating-point effect depends on it (Section 3.11).

(ii) The study covers one household, and the outages are a modelled timetable, not recorded interruptions; October to April have no scheduled outages, which understates generator use.

(iii) The PV series is the production measured at one plant and shows little day-to-day variability, so PV forecasting is easy and PV variability is probably understated; the no-storage case cannot use PV during an outage and curtails PV surplus, which favours storage in the primary-energy balance.

(iv) The perfect-foresight schedule is an upper bound, not an achievable result.

(v) Battery degradation, temperature effects on the generator and PV and generator start-up fuel are not modelled.

(vi) The grid emission factor is a reported national average [35] and the primary-energy factor is a derived estimate; marginal factors during outage hours may differ.

(vii) The scenario model has three points per step and treats the forecast noise as independent between steps; the error quantiles are estimated from one household over one year, with a coverage of 78.5% against 80%; the comparison of forecasters and its link to fuel use were run in a Python replication over the outage season, and only the ridge-based controllers were run in the main (MATLAB) simulation over the full year.

5. Conclusions

For a Baghdad-type supply timetable, a 5 kWh/5 kW household battery reduces diesel use from about 2290 L per year to about 212 L with a forecast-aware scenario controller, a reduction of about 91% against no storage and 82% against rule-based storage; the saving exceeds the energy displaced because the remaining generator operation takes place at about 25% electrical efficiency instead of about 12%. Grid pre-charge provides most of the saving; choosing the generator operating point adds about a quarter of the remaining diesel at 2.5–5 kWh, an effect that depends on the idle fuel of the generator. The scenario controller adds a further 20.5% (5 kWh) and 18.5% (2.5 kWh) against grid pre-charge, 2.3–3.7% more than a point-forecast controller, and it realises 76–86% of the gain of a perfect-foresight schedule. A multilayer perceptron, an LSTM network and gradient boosting did not improve fuel use over a ridge-corrected seasonal forecast, even where they were more accurate, so the useful step was to plan around the forecast error. The net primary-energy balance remains positive when grid pre-charge is counted (the fuel energy saved is 23 times the extra grid energy at 5 kWh), and the scenario controller avoids 5.1 t of CO_2 per household per year at the reported Iraqi grid factor, against 5.6 t for the perfect-foresight schedule at

zero grid factor. Its monetary advantage over a point-forecast controller is small (about 1–3 $ per year); the benefit is the reduction of fuel and generator hours. The extra benefit of choosing the generator operating point depends on the idle fuel (17–32% at 5 kWh for the perfect-foresight schedule), which must be measured. These results rest on an assumed fuel curve and a modelled timetable; measuring the generator's fuel use, adding recorded outage data and extending the analysis to several households are the next steps.

Author Contributions: Conceptualization, A.M. and B.A.; methodology, A.M., Q.Z. and A.R.; software, A.M. and N.A.; validation, A.M., O.A. and N.A.; formal analysis, A.M. and O.A.; investigation, A.M.; resources, J.C. and A.R.; data curation, A.M. and N.A.; writing,original draft preparation, A.M.; writing,review and editing, A.M., B.A., A.R., Q.Z., O.A. and J.C., with language editing by N.A.; visualization, A.M. and O.A.; supervision, B.A., J.C. and A.R.; project administration, A.R. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

A, B, C	No storage; rule-based storage; storage with grid pre-charge
BTE	Brake thermal efficiency
CRF	Capital recovery factor
DP	Dynamic programming
IQD	Iraqi dinar
LHV	Lower heating value
LSTM	Long short-term memory
MAE	Mean absolute error
MLP	Multilayer perceptron
MPC	Model predictive control
MPC1	Point-forecast model predictive controller
OPT	Perfect-foresight optimal schedule
PE	Primary energy
PV	Photovoltaic
RMSE	Root-mean-square error
SFC	Specific fuel consumption
SMPC	Scenario-based (stochastic) model predictive controller
SOC	State of charge
SS	Skill score

References

1. 964media. Baghdad Residents Face Scheduled Power Cuts and Rely on Private Generators. Available online: https://en.964media.com/47717/ (accessed on 22 October 2026).
2. Al-Kayiem, H.H.; Mohammad, S.T. Potential of Renewable Energy Resources with an Emphasis on Solar Power in Iraq: An Outlook. *Resources* **2019**, *8*, 42. [CrossRef]
3. Al-Bayati, A.M.S.; Abdulkhudhur, H.H. Optimal Sizing and Analysis of a Hybrid Solar-Wind-Diesel-Battery Energy System Considering Distinct Tracking Systems in Iraq. *Int. J. Energy Prod. Manag.* **2026**, *11*, 323–345. [CrossRef]
4. Iraq Business News. Iraq Cuts Diesel Price for Private Generator Operators. 26 May 2026. Available online: https://www.iraq-businessnews.com/2026/05/26/iraq-cuts-diesel-price-for-private-generator-operators/ (accessed on 22 October 2026).
5. Dwisyahputra, B.E.; Setiawan, E.A. Transforming Energy Backup: Evaluating Battery Storage vs. Generators for Power Outages in Rooftop Solar Systems in Residential Buildings of Indonesia's New Capital City. *SSRN* **2024**. [CrossRef]
6. Ahmad, I.; Ahmad, S.; Arif, M.Z. Optimizing lead-acid battery sizing for household UPS systems under variable power outages. *Next Res.* **2025**, *2*, 101023. [CrossRef]

7. Thirumalai, S.K.; Karthick, A.; Dhal, P.K.; Pundir, S. Photovoltaic-wind-battery and diesel generator-based hybrid energy system for residential buildings in smart city Coimbatore. *Environ. Sci. Pollut. Res.* **2024**, *31*, 14229–14238. [CrossRef] [PubMed]
8. Silva, C.M.M.R.S.; Premadasa, P.N.D.; Chandima, D.P.; Karunadasa, J.P.; Wheeler, P. Optimum sizing and configuration of electrical system for telecommunication base stations with grid power, Li-ion battery bank, diesel generator and solar PV. *J. Energy Storage* **2025**, *123*, 116704. [CrossRef]
9. Pillai, K.; Sundaram, S. Optimization and Feasibility Analysis of Hybrid Distributed Generator Based System with a Comparison of Battery and Hydrogen Energy Storage for Residential Electrification. *Energy Storage* **2024**, *6*, e70075. [CrossRef]
10. Heywood, J.B. *Internal Combustion Engine Fundamentals*, 2nd ed.; McGraw-Hill Education: New York, NY, USA, 2018.
11. Guzzella, L.; Onder, C.H. *Introduction to Modeling and Control of Internal Combustion Engine Systems*; Springer: Berlin/Heidelberg, Germany, 2010. [CrossRef]
12. DEUTZ AG. F/FL 912 Genset Engines, Datasheet. Available online: https://www.deutz.com/fileadmin/contents/global/produkte/datasheets/en/legacy/FL_912_Genset_EN.pdf (accessed on 22 October 2026).
13. Lambert, T.; Gilman, P.; Lilienthal, P. Micropower System Modeling with HOMER. In *Integration of Alternative Sources of Energy*; Wiley: Hoboken, NJ, USA, 2005; pp. 379–418. [CrossRef]
14. Dufo-López, R.; Bernal-Agustín, J.L. Multi-objective design of PV–wind–diesel–hydrogen–battery systems. *Renew. Energy* **2008**, *33*, 2559–2572. [CrossRef]
15. Ghanbari, K.; Maleki, A. Optimal design of off-grid hybrid PV–wind–battery–diesel systems under component degradation using a metaheuristic algorithm. *Energy Convers. Manag.* **2026**, *348*, 120652. [CrossRef]
16. Parisio, A.; Rikos, E.; Glielmo, L. A Model Predictive Control Approach to Microgrid Operation Optimization. *IEEE Trans. Control Syst. Technol.* **2014**, *22*, 1813–1827. [CrossRef]
17. Zia, M.F.; Elbouchikhi, E.; Benbouzid, M. Microgrids energy management systems: A critical review on methods, solutions, and prospects. *Appl. Energy* **2018**, *222*, 1033–1055. [CrossRef]
18. Hirsch, A.; Parag, Y.; Guerrero, J. Microgrids: A review of technologies, key drivers, and outstanding issues. *Renew. Sustain. Energy Rev.* **2018**, *90*, 402–411. [CrossRef]
19. Ahmed, F.; Al-Abri, R.; Yousef, H.; Massoud, A.M. An Optimal Energy Dispatch Management System for Hybrid Power Plants: PV-Grid-Battery-Diesel Generator-Pumped Hydro Storage. *IEEE Access* **2024**, *12*, 143307–143326. [CrossRef]
20. Mohamed, M.M.; El Zoghby, H.M.; Sharaf, S.M.; Mosa, M.A. Optimal virtual synchronous generator control of battery/supercapacitor hybrid energy storage system for frequency response enhancement of photovoltaic/diesel microgrid. *J. Energy Storage* **2022**, *51*, 104317. [CrossRef]
21. Alarcón, M.A.; Alarcón, R.G.; González, A.H.; Ferramosca, A. A scenario-based economic-stochastic model predictive control for the management of microgrids. *Sustain. Energy Grids Netw.* **2023**, *36*, 101205. [CrossRef]
22. Abdelghany, M.B.; Al-Durra, A.; Zeineldin, H.; Gao, F. Integrating scenario-based stochastic-model predictive control and load forecasting for energy management of grid-connected hybrid energy storage systems. *Int. J. Hydrogen Energy* **2023**, *48*, 35624–35638. [CrossRef]
23. Babić, L.; Lauricella, M.; Ceusters, G.; Biskoping, M. Data-driven non-parametric chance-constrained model predictive control for microgrids energy management using small data batches. *Front. Control Eng.* **2023**, *4*, 1237759. [CrossRef]
24. Xu, B.; Oudalov, A.; Ulbig, A.; Andersson, G.; Kirschen, D.S. Modeling of Lithium-Ion Battery Degradation for Cell Life Assessment. *IEEE Trans. Smart Grid* **2018**, *9*, 1131–1140. [CrossRef]
25. Vykhodtsev, A.V.; Jang, D.; Wang, Q.; Rosehart, W.; Zareipour, H. Physics-Aware Degradation Model of Lithium-ion Battery Energy Storage for Techno-Economic Studies in Power Systems. *IEEE Trans. Sustain. Energy* **2024**, *15*, 316–327. [CrossRef]
26. Stanchev, P.; Hinov, N. Comparative Techno-Economic and Life Cycle Assessment of Stationary Energy Storage Systems: Lithium-Ion, Lead-Acid, and Hydrogen. *Batteries* **2025**, *11*, 382. [CrossRef]
27. Ntube, N.; Li, H. Stochastic multi-objective optimal sizing of battery energy storage system for a residential home. *J. Energy Storage* **2023**, *59*, 106403. [CrossRef]
28. Barley, C.D.; Winn, C.B.; Flowers, L.; Green, H.J. *Optimal Control of Remote Hybrid Power Systems. Part 1: Simplified Model*; Technical Report NREL/TP-441-7806; National Renewable Energy Laboratory: Golden, CO, USA, 1995. Available online: https://www.osti.gov/biblio/52813 (accessed on 19 November 2026).
29. Aziz, A.S.; Tajuddin, M.F.N.; Adzman, M.R.; Ramli, M.A.M.; Mekhilef, S. Energy Management and Optimization of a PV/Diesel/Battery Hybrid Energy System Using a Combined Dispatch Strategy. *Sustainability* **2019**, *11*, 683. [CrossRef]
30. Mohammed, A.; Abdullah, B.M.; Shubbar, A.; Zhang, Q.; Aldhaibani, O.; Cullen, J.; Salih, A. Deep Reinforcement Learning for Battery Energy Storage Optimization and Residential Decarbonization in Grid-Deficient Environments: An Iraqi Case Study. *Energies* **2026**, *19*, 1233. [CrossRef]
31. HOMER Energy. Generator Fuel Curve Intercept Coefficient. HOMER Pro 3.15 Documentation. Available online: https://homerenergy.com/products/pro/docs/3.15/generator_fuel_curve_intercept_coefficient.html (accessed on 22 October 2026).

32. Enviraj. Diesel Generator Fuel Consumption Chart (Approximate). 2022. Available online: https://oer.enviraj.com/general/diesel-generator-fuel-consumption-chart/ (accessed on 22 October 2026).
33. Ministry of Electricity of Iraq. How the Electricity Bill Is Calculated (Tariff Examples). 5 July 2022. Available online: https://www.moelc.gov.iq (accessed on 22 October 2026).
34. U.S. Energy Information Administration. Iraq: Country Analysis Brief. 2025. Available online: https://www.eia.gov/international/analysis/country/IRQ (accessed on 22 October 2026).
35. Low Carbon Power. Carbon Intensity of Electricity in Iraq. Available online: https://lowcarbonpower.org/region/Iraq (accessed on 22 October 2026).
36. Rennert, K.; Errickson, F.; Prest, B.C.; Rennels, L.; Newell, R.G.; Pizer, W.; Kingdon, C.; Wingenroth, J.; Cooke, R.; Parthum, B.; et al. Comprehensive evidence implies a higher social cost of CO_2. *Nature* **2022**, *610*, 687–692. [CrossRef] [PubMed]
37. Kesicki, F.; Strachan, N. Marginal abatement cost (MAC) curves: Confronting theory and practice. *Environ. Sci. Policy* **2011**, *14*, 1195–1204. [CrossRef]
38. DailyFuels. Iraq Fuel Prices. Available online: https://www.dailyfuels.com/iraq/ (accessed on 22 October 2026).
39. Energy-Storage.news. Battery Storage System Prices Continue to Fall Sharply, BNEF and Ember Reports Find. 2025. Available online: https://www.energy-storage.news/?p=92133 (accessed on 22 October 2026).
40. National Renewable Energy Laboratory. Residential Battery Storage. In *Annual Technology Baseline 2024*; NREL: Golden, CO, USA, 2024. Available online: https://www.osti.gov/biblio/2377191 (accessed on 19 November 2026).
41. Kong, W.; Dong, Z.Y.; Jia, Y.; Hill, D.J.; Xu, Y.; Zhang, Y. Short-Term Residential Load Forecasting Based on LSTM Recurrent Neural Network. *IEEE Trans. Smart Grid* **2019**, *10*, 841–851. [CrossRef]
42. Haben, S.; Arora, S.; Giasemidis, G.; Voss, M.; Vukadinović Greetham, D. Review of low voltage load forecasting: Methods, applications, and recommendations. *Appl. Energy* **2021**, *304*, 117798. [CrossRef]
43. International Finance Corporation. *The Dirty Footprint of the Broken Grid: The Impacts of Fossil Fuel Back-Up Generators in Developing Countries*; IFC: Washington, DC, USA, 2019.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.