Svedberg Open
DISSEMINATION OF KNOWLEDGE

International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: https://www.svedbergopen.com/



Research Paper Open Access

Value of Forecast Accuracy and Horizon in Predictive Energy Management of An Islanded Saharan PV–Wind–Battery Microgrid With A Deferrable Water Load: A Nine-Year Simulation Study

Salim Kadri [1], Djalal Djarah [1,*], Ahmed Djafour [1]

[1] Faculty of Applied Sciences, LAGE Laboratory, Kasdi Merbah University of Ouargla, Ouargla 30000, Algeria.
[*] Corresponding author: Djalal Djarah, Faculty of Applied Sciences, LAGE Laboratory, Kasdi Merbah University of Ouargla, Ouargla 30000, Algeria; e-mail: djarah.djalal@univ-ouargla.dz
E-mail addresses: kadri.salim@univ-ouargla.dz (S. Kadri); djarah.djalal@univ-ouargla.dz (D. Djarah); djafour.ahmed@univ-ouargla.dz (A. Djafour)

Received: 02/08/2026 ; Accepted: 20/09/2026

Abstract
In an islanded renewable microgrid without a backup generator, control quality depends on the information available about the coming hours. We quantify how forecast accuracy and horizon affect the management of a Saharan PV–wind–battery system (Ouargla). The modelled load comprises a critical part (65%) and a deferrable water station (35%) with a tank. We use nine years of hourly NASA POWER weather (2016–2024) and simulated, lead-time-dependent forecast errors to compare four strategies: a perfect-information oracle, a linear model predictive controller (MPC; 48 h horizon), a rule tuned on an independent year (2015) and a strategy without flexibility. The MPC criterion exceeds that of the oracle by 8.5% (95% bootstrap interval over years: 7.7–9.4%), against 53.4% for the tuned rule and 242% without flexibility. Unserved critical energy is 0.54, 3.33 and 11.46 MWh/year, respectively. Five additional error draws give nine-year mean excesses of 8.7–9.3%. With perfect forecasts the excess is 0.5%, and horizons shorter than 24 h degrade performance in this configuration. With a naive 24 h persistence forecast the excess rises to 59%, no better than the tuned rule, although the critical load stays protected. The results quantify performance under the stated load, component and forecast-error assumptions, and rest entirely on simulation. The fair-reference protocol (tuning on an independent year, paired nine-year evaluation) can be reused to assess other controllers.

Key words: model predictive control; islanded microgrid; deferrable load; lithium-ion battery; value of forecasting; Sahara.

1. Introduction

Isolated Saharan sites have one of the highest solar resources in the world. Their reliability, however, still usually rests on a diesel generator and a fixed operating rule whose distance from the optimum is rarely quantified. This article studies a fully renewable microgrid, with neither diesel nor hydrogen (PV, wind and a lithium-ion battery), in which part of the load is deferrable: the water supply, provided by a pumping or treatment station and a tank. The tank acts as a storage without conversion losses but with a service constraint. Energy management must therefore trade off the reliability of the critical load, the satisfaction of water demand and the ageing of the battery. Economic model predictive control is the usual answer to this kind of trade-off (Hu et al., 2024). It has been validated on a real case for an autonomous microgrid with shiftable loads (Cecilia et al., 2020) and compared with rule-based variants (Pippia et al., 2020). Robust and stochastic formulations extend it to uncertainty (Zhang et al., 2018), and it has been applied to the coordination of desalination plants with renewable-rich power systems (Elsir et al., 2023). The influence of the ageing model on management is known (Wang et al., 2020; Wicke & Bocklisch, 2024). The Algerian studies of autonomous microgrids that we found mostly address sizing, including with demand management (Boutaghane et al., 2025; Mokhtara et al., 2020), while the performance of Saharan PV plants is documented separately (Chaib et al., 2024). What is missing is an evaluation of the real contribution of prediction against a properly tuned rule, over several years of real weather, with a flexible load and forecast errors that grow with lead time. A controller compared with an untuned rule yields gains that depend on the tuning of the reference as much as on the controller.

Recent work on islanded and isolated systems bears directly on this question. Manzano et al. (2022) ran economic predictive control on real demand and renewable data with forecast errors and reported savings of more than 10% over non-predictive methods. Gholami and Pisano (2022) showed that including the conversion losses of storage in the predictive model limits errors in the predicted state of charge. The cost of poor forecasts is documented for a grid-connected energy community: dispatch by predictive control saved 24.7% of operating cost with perfect foresight but only 3.3% with real forecasts (Houben et al., 2023). Battery wear also matters. Leaving degradation out of the schedule understates the cost of storage (Mazidi et al., 2023), and simple current derating based on state of charge and temperature was estimated to lengthen battery life by about 45% in operating mini-grids (Sowe et al., 2022). On the water side, an Algerian off-grid PV–wind–hydrogen study treats water extraction and treatment as a deferrable load (Boutaghane et al., 2025). For the Saharan resource itself, measured data from a 3 MWp plant in the Algerian desert give an annual performance ratio of 67% (Chaib et al., 2024), so modelled PV output should be checked against field data. None of these works isolates, for an islanded site with a deferrable water load, how much of the gap to a perfect-information optimum comes from the forecast, the horizon and the tuning of the reference rule.

Our contributions are as follows.

1. A linear formulation of the management of a PV–wind–battery microgrid with a deferrable water load, solved in a few milliseconds at every hour and completed by a real-time balancing layer.
2. A fair-reference evaluation protocol: a rule tuned on an independent year, then nine years of real weather (NASA POWER, Ouargla) as evaluation years, with paired differences.
3. A quantification of the value of flexibility (tank size, flexible share), of the value of prediction (forecast error, horizon, terminal values) and of the comparison with a rule-based predictive controller from the literature (Pippia et al., 2020).
4. An analysis of ageing by rainflow counting and a calendar law, and of what a state-of-charge holding cost changes in reliability and lifetime.

The extension to a diesel generator and a hydrogen chain is left to later work. Section 2 describes the system, the data, the control problem, the compared strategies and the protocol. Section 3 presents the results and their discussion, and Section 4 the conclusion.

2. Materials and Methods

2.1 System architecture

The microgrid is islanded and treated as an equivalent DC system at an hourly time step. An 800 kWp PV array and a 150 kW wind turbine supply two loads: a critical load (65% of annual energy), which must be served every hour, and a water load (35%), deferrable within the limits of a tank. A 2.5 MWh, 500 kW lithium-ion battery provides electrical storage. The set is not sized jointly with the controller; the sizes are data of the problem (Table 1).

The water load is represented as an equivalent electrical power. A pumping or treatment station modulates its power between 0 and 150 kW and fills a tank whose capacity equals 12 h of mean demand (574 kWh electrical equivalent). The water demand is drawn from this tank, and demand not satisfied when the tank is empty is a water shortage. The tank has no energy losses. The continuous modulation of the station power and the absence of a minimum power are simplifications, discussed in Section 3.7. The architecture is shown in Figure 1.

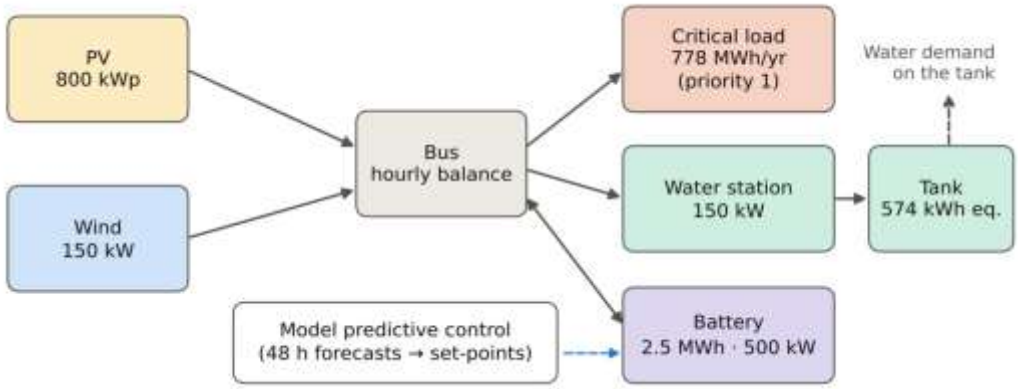


Figure 1. Architecture of the microgrid and of the control: the hourly balance supplies the critical load (priority 1) and the water station, which fills the tank; the battery exchanges with the bus; predictive control sets the set-points from 48 h forecasts.

Table 1. System parameters.

Component	Parameter	Value
PV	Peak power	800 kWp
Wind	Rated power	150 kW (hub height 30 m)
Battery	Energy; power	2.5 MWh; 500 kW
Battery	Round-trip efficiency; state-of-charge range	0.90; 15–95%
Water station	Maximum power	150 kW
Tank	Capacity	574 kWh electrical eq. (12 h of mean demand)
Critical load	Mean; peak; annual energy	88.8 kW; 150.9 kW; 778 MWh
Water load	Mean; peak; annual energy	47.8 kW; 84.6 kW; 419 MWh

2.2 Real weather data

Weather comes from NASA POWER for Ouargla (31.95° N, 5.33° E), hourly series 2015–2024 (NASA POWER Project, 2026): global, direct and diffuse horizontal irradiance, 2 m temperature and 10 m wind speed. The 29 February of leap years (2016, 2020, 2024) is removed to obtain 8,760 values per year. Annual global horizontal irradiation varies from 2,091 to 2,168 kWh/m^2.

PV power is computed in the plane of an array tilted by $\beta = 28°$ due south, at latitude $\varphi = 31.95°$ N. For day of year n and hour h, the solar declination δ, the hour angle ω (solar time at mid-hour) and the cosine of the angle of incidence on the tilted south-facing plane are

$$\delta(n) = 23.45° \cdot \sin[2\pi (284 + n) / 365] \quad (1)$$

$$\omega(h) = 15° \cdot (h + 0.5 - 12) \quad (2)$$

$$\cos\theta = \sin\delta \sin\varphi \cos\beta - \sin\delta \cos\varphi \sin\beta + \cos\delta \cos\varphi \cos\beta \cos\omega + \cos\delta \sin\varphi \sin\beta \cos\omega \quad (3)$$

Equation (1) is the usual sinusoidal approximation of the declination, and equations (2) and (3) are the geometric relations of solar geometry for a south-facing surface. With the isotropic transposition of the sky diffuse irradiance and a ground albedo $\rho g = 0.27$, the plane-of-array irradiance is

$$GPOA = DNI \cdot \max(\cos\theta, 0) \cdot 1[\cos\theta z > 0.02] + DHI\,(1 + \cos\beta)/2 + GHI\,\rho g\,(1 - \cos\beta)/2 \quad (4)$$

where DNI, DHI and GHI are the direct normal, diffuse horizontal and global horizontal irradiances of the NASA POWER series, and $\cos\theta z = \sin\varphi \sin\delta + \cos\varphi \cos\delta \cos\omega$ is the cosine of the solar zenith angle. Soiling is represented by a sawtooth: the loss sd grows by 0.004 per dry day, is capped at 0.15 and is reset to 0.01 every 30 days (no rain washing, for lack of precipitation data),

$$sd = \min(sd_{-1} + 0.004;\ 0.15);\ sd \leftarrow 0.01 \text{ if } d \bmod 30 = 29 \quad (5)$$

Soiling accumulates during dry periods and is removed by cleaning or rain, an effect that is documented (Kimber et al., 2006); the numerical rate is an assumption for the site. The cell temperature follows the linear correlation with irradiance, and the PV power a linear dependence on cell temperature, two forms reviewed by Skoplaki and Palyvos (2009):

$$Tc = Ta + 0.034\,GPOA \quad (6)$$

$$PPV = PPV_{nom} \cdot (GPOA/1000) \cdot [1 - 0.0038\,(Tc - 25)] \cdot 0.92 \cdot (1 - sd) \quad (7)$$

with Ta the 2 m air temperature (°C), 0.0038/K the power temperature coefficient and 0.92 the system efficiency. The power is limited to 0.95 PPV_{nom}. The PV capacity factor is 19.1–20.2% depending on the year. For wind, the speed measured at 10 m is brought to the hub height of 30 m by a power law, then converted by a cubic curve between the cut-in speed vci = 3 m/s, the rated speed vr = 11 m/s and the cut-out speed vco = 25 m/s:

$$v30 = v10\,(30/10)^{\alpha},\ \alpha = 1/7 \quad (8)$$

$$PWT = PWT_{nom} \cdot (v30^3 - vci^3) / (vr^3 - vci^3) \text{ for } vci \le v30 < vr;\ PWT_{nom} \text{ for } vr \le v30 < vco;\ 0 \text{ otherwise} \quad (9)$$

The exponent 1/7 is a common default value, not a site measurement. It depends on atmospheric stability and surface roughness, so the wind capacity factor, 11.2–13.7%, depends on this height assumption. Renewable energy represents 124–133% of the load energy depending on the year (Figure 2).

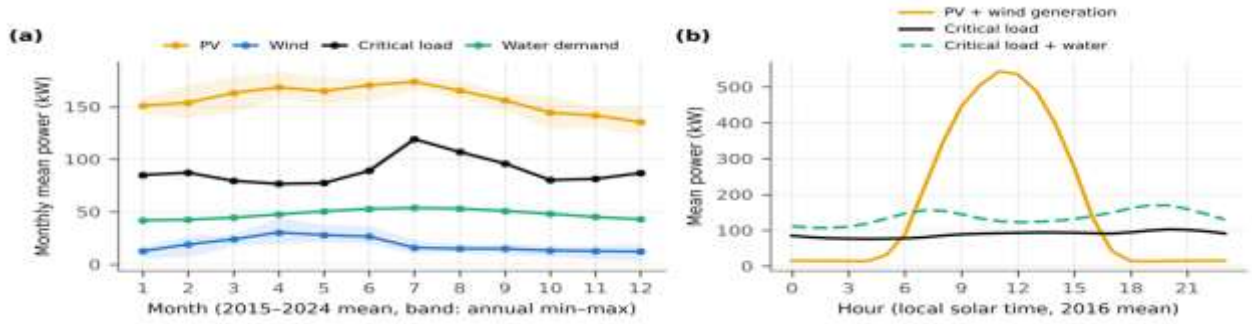


Figure 2. (a) Monthly mean PV, wind, critical and water powers (2015–2024 mean, band: annual minimum–maximum). (b) Mean daily profile of production and loads in 2016.

2.3 Loads

The total load is the published hourly shape of 2018 of the "Load Consumption Data Algeria" dataset (Load Consumption Data Algeria, 2023), scaled to 1,197 MWh/year (mean 136.6 kW). It is an aggregated national shape, not a record of the site. The critical load is 65% of it, with the same shape. Water demand is a synthetic profile with two peaks (7 h and 19 h) and a seasonal variation of ±12.5% (maximum in July), normalised to 35% of the total energy. It is assumed to be known to the controller (scheduled service). These two profiles are modelling assumptions, so the result is established for this load structure only; Section 3.4 varies the share and the tank.

2.4 Battery and ageing

Cycling ageing is computed from the state-of-charge trajectory by 3-point rainflow counting (Downing & Socie, 1982). This gives cycles i of amplitude di (fraction of nominal capacity) and count ni (1 for a full cycle, 0.5 for a half cycle). A Wöhler law gives the number of cycles to failure, and Miner's linear accumulation gives the annual cycling damage:

$$N(d) = 6000 \cdot (0.8 / d)^{1.7} \quad (10)$$

$$Dcyc = \Sigma i\ ni / N(di) \quad (11)$$

Calendar ageing is a function of the mean state of charge over the year; the dependence of calendar ageing on the storage state of charge is documented (Keil et al., 2016). The exponential form and its constants are modelling assumptions:

$$Dcal = (1/15\ yr) \cdot \exp[1.5\ (SOCmean - 0.5)] \quad (12)$$

$$Tlife = 1 / (Dcyc + Dcal) \quad (13)$$

The calendar life is thus 15 years at a mean state of charge of 50%. The Wöhler parameters (6,000 cycles at 80% amplitude, exponent 1.7) and the linear damage accumulation are assumptions of the model, not measurements on a specific cell. Temperature, which would accelerate ageing in the Sahara, is not modelled.

2.5 Forecast errors

At every hour the controller receives 48 h forecasts of PV power, wind power and critical load. The errors are multiplicative and correlated across lead time (AR(1), $\rho = 0.9$). Their standard deviation grows with lead time from 3% to 25% for PV, from 8% to 40% for wind and from 2% to 7% for the critical load. For a quantity x with true value x(t+j) at lead j = 0, …, H − 1, the forecast is

$$\hat{x}(t+j) = \max\{0;\ x(t+j)\ [1 + \sigma x(j)\ zj]\},\ z0 \sim N(0,1),\ zj = \rho\ zj{-}1 + \sqrt{(1 - \rho^2)}\ \varepsilon j,\ \rho = 0.9 \quad (14)$$

with εj independent standard normal variables, and $\sigma x(j) = \sigma 0 + (\sigma\infty - \sigma 0)\ [1 - \exp(-j/\tau)]$ ($\sigma 0 = 0.03$, $\sigma\infty = 0.25$, $\tau = 14$ h for PV; 0.08, 0.40, 12 h for wind; 0.02, 0.07, 24 h for the load, respectively). The errors are simulated and are not produced by a real weather-forecast model. An operational reference without any simulated error is also used in Section 3.3: a 24 h persistence forecast, in which the forecast of each hour is the value observed at the same hour of the previous available day, for PV, wind and load.

2.6 Management problem

At every hour k, the controller (receding-horizon, i.e. model predictive control; Mayne et al., 2000) chooses the battery charge power Pch and discharge power Pdis, the water station power Ppump, the unserved critical power Pun, the curtailed power Pcurt and the water shortage Sw, all non-negative. With the renewable power PRE, the critical load Lc and the water demand Dw, the power balance and the states of the battery (energy Eb) and of the tank (content W) are

$$PRE(k) + Pdis(k) + Pun(k) = Lc(k) + Pch(k) + Ppump(k) + Pcurt(k) \quad (15)$$

$$Eb(k+1) = Eb(k) + \eta b\ Pch(k)\ \Delta t - Pdis(k)\ \Delta t / \eta b \quad (16)$$

$$W(k+1) = W(k) + [Ppump(k) - Dw(k) + Sw(k)]\ \Delta t \quad (17)$$

with $0.15\ Eb_{nom} \le Eb \le 0.95\ Eb_{nom}$, $0 \le W \le Wmax$, $Pch, Pdis \le 500$ kW, $Ppump \le 150$ kW, $Pun \le Lc$ and $Sw \le Dw$. Over a horizon of N steps, the controller minimises

$$J = \Sigma k{=}1..N\ [cun\ Pun + cw\ Sw + cdeg\ Pdis + ch\ Eb]\ \Delta t - vb\ Eb(N) - vw\ W(N) \quad (18)$$

All coefficients are dimensionless weights (p.u./kWh) with no monetary value: cun = 3 (unserved critical load), cw = 1 (water shortage), cdeg = 0.02 (wear proportional to the discharged energy), ch ≥ 0 (state-of-charge holding, zero except in Section 3.5), and vb = 0.15 and vw = 0.30 (terminal values). The ratio cun/cw expresses that critical load is served before water. Charging and then discharging dissipates energy ($\eta b^2 < 1$) and costs cdeg, so simultaneous charging and discharging is never optimal. The problem is therefore a linear program without binary variables, solved with the HiGHS solver (Huangfu & Hall, 2018). The terminal terms vb Eb(N) and vw W(N) play the role of a terminal cost, which economic model predictive control uses to account for what remains stored at the end of the horizon (Amrit et al., 2011; Mueller et al., 2014).

2.7 Compared strategies

NF, no flexibility. Water demand is a non-deferrable load served at the hour it occurs; in case of deficit, water is shed before the critical load. It is the reference for the value of flexibility.

S1, default rule. Surplus charges the battery up to 95%, then feeds the water station, then is curtailed. A deficit is covered by the battery down to 15%, and the battery never feeds the water station.

S2, tuned rule. Same logic with four parameters: a state-of-charge threshold above which surplus goes first to the water station (one value in the morning, another after 13:00), the tank level below which the battery may feed the station, and the minimum state of charge for doing so. The values are chosen by exhaustive search (192 combinations) on a calibration year, 2015, which is not used for evaluation. S2 is the fair reference: it represents what a careful engineer obtains without prediction.
S3, oracle. The linear program over the whole year (8,760 steps), with actual production and load. It is not implementable and bounds J from below.
S4, flat MPC. At every hour, the controller receives the forecasts over H = 48 h, solves the linear program (HiGHS), applies the first set-points and starts again. Terminal values are constant.
R, rule-based MPC (Pippia et al., 2020). Same program, but the mode is fixed by the sign of the forecast net balance PRE − Lc: with a forecast surplus discharge is forbidden, and with a forecast deficit charging and pumping are forbidden.
S5, ageing-aware MPC. S4 with ch > 0, a linearisation of the calendar cost that grows with the mean state of charge.
The real-time balancing layer, common to the predictive strategies, applies the first-hour set-points to the actual production and load. In an actual surplus it first cancels the planned discharge, then charges the battery, then feeds the water station, then curtails. In an actual deficit it first reduces the water station, then the planned charge, then discharges the battery further, and the remainder is unserved critical load.

2.8 Experimental protocol

Rule S2 is tuned on 2015. The nine years 2016–2024 are evaluation years: each strategy is simulated there over 8,760 hours, with the battery and tank at 50% at the start. The forecast errors of each year are a single draw, with a seed equal to the year, identical for all predictive strategies. Comparisons are paired by year. The final states of the battery and the tank are not constrained, criterion J excludes terminal values, and annual boundary effects were not analysed separately.
The main criterion is the criterion J of Section 2.6 (without terminal values) relative to the oracle of the same year. It is completed by unserved critical energy (in MWh/year and as a percentage of the critical load), water shortage (percentage of demand), curtailed energy (percentage of renewable energy), mean battery state of charge, battery lifetime (Section 2.4) and mean solving time per hour.
Means and standard deviations are taken over the evaluation years. With nine years, we do not run hypothesis tests; we report instead the number of years in which each paired difference is positive. The standard deviation mixes weather variability and one forecast-error draw per year, and Section 3.1 checks the effect of the draw with five additional ones. The sensitivity analyses (Sections 3.3 to 3.6) use three years, 2016, 2020 and 2024, to bound computing time, and S2 is re-tuned on 2015 for each configuration.

3. Results and Discussion

3.1 Value of prediction over nine real years

Over the nine evaluation years, the flat predictive controller (S4) is at 8.5 ± 1.4% of the oracle for criterion J, against 53.4 ± 3.1% for the tuned rule (S2). S4 reduces J by 29.2% and recovers 84% of the gap between S2 and the oracle, and the difference has the same sign in each of the nine years (Table 2, Figure 3a).

Table 2. Evaluation over nine years of real data (2016–2024). Mean ± interannual standard deviation; J in thousands of p.u.; gap: excess of J over the oracle of the same year; SOC: mean state of charge; lifetime: Section 2.4.

Strategy	J (10^3 p.u.)	Gap (%)	Unserved critical load (MWh/yr)	Water shortage (%)	Curtailment (% RE)	Mean SOC (%)	Lifetime (years)
S3 Oracle	17.9 ± 4.8	0	0.19 ± 0.33	2.26 ± 1.01	20.1	37	18.0 ± 0.2
S4 MPC, perfect forecast	18.0 ± 4.9	0.5 ± 0.3	0.19 ± 0.33	2.28 ± 1.03	20.0	70	11.0 ± 0.2
S4 flat MPC	19.4 ± 5.2	8.5 ± 1.4	0.54 ± 0.41	2.33 ± 1.00	20.0	70	10.8 ± 0.1
S5 ageing-aware MPC (ch = 0.003)	22.7 ± 5.7	27.8 ± 4.0	1.51 ± 0.65	2.09 ± 0.92	19.5	65	11.6 ± 0.2
S2 Tuned rule	27.5 ± 7.8	53.4 ± 3.1	3.33 ± 1.61	2.06 ± 0.77	19.8	67	11.6 ± 0.2

R Rule-based MPC	56.2 ± 3.5	230 ± 67	0.24 ± 0.35	11.54 ± 0.69	22.8	75	10.3 ± 0.1
NF No flexibility	59.9 ± 10.8	242 ± 33	11.46 ± 2.91	3.32 ± 0.55	19.7	59	12.7 ± 0.1
S1 Default rule	146.3 ± 5.3	765 ± 200	0.19 ± 0.33	33.08 ± 1.18	28.7	77	9.9 ± 0.0

To check that the result does not depend on the particular error draw, S4 was rerun over the nine years with five additional independent draws (seeds 101–105, 45 runs). The mean excess over the oracle is 9.0% (range of the nine-year mean over the five draws: 8.7–9.3%, standard deviation 0.25 points). The standard deviation across draws within a year is 0.8 points on average (0.6–1.4 points depending on the year). Over the six draws of each year, the mean excess is 8.9%. S4 gives a lower J than the tuned rule S2 in all 45 additional runs (maximum excess 11.8%, against 53.4% for S2). The draw of the main results (8.5%) therefore lies slightly below the average of the six draws, and the ordering of the strategies does not depend on it. Resampling the nine years with replacement (10,000 draws, paired by year) gives 95% intervals of 7.7–9.4% for the excess of S4 (main draw; 8.0–9.9% for the six-draw mean), 51.5–55.4% for S2 and 222–263% for NF. The corresponding intervals are 27.9–30.5% for the reduction of J relative to S2, 2.1–3.6 MWh/year for the unserved critical load avoided relative to S2, and 0.76–0.87 years for the lifetime difference. With nine years, these intervals describe weather variability under the assumed models, not uncertainty about a real site.

Unserved critical load is 0.54 ± 0.41 MWh/year for S4 (0.07% of the critical load), against 3.33 ± 1.61 MWh/year for S2 (0.43%), i.e. 2.8 MWh/year less in each of the nine years. Without flexibility (NF) it reaches 11.46 ± 2.91 MWh/year (1.47%), 21 times more than S4. Reserving part of the load for a buffer tank and driving it with forecasts thus reduces J by 68%. Water shortage is similar for S4 (2.33%) and S2 (2.06%). S2 serves water slightly better, but at the expense of the critical load, a loss that criterion J, which weighs the critical load three times more, does not offset. Curtailment, close to 20% of renewable energy, is structural, since renewable energy represents 124–133% of the load.

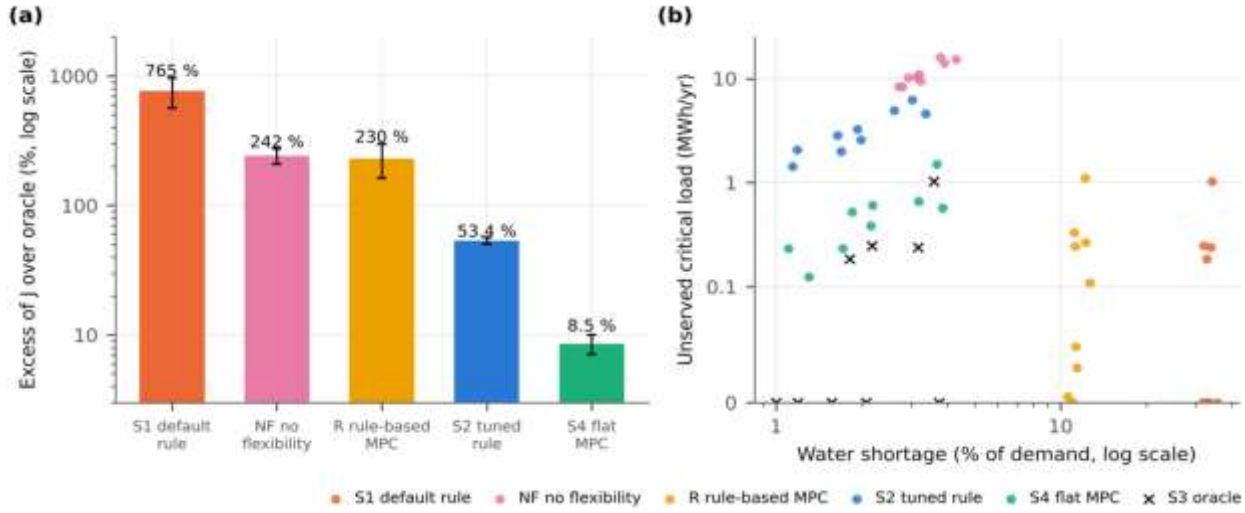


Figure 3. (a) Excess of criterion J over the oracle, mean and standard deviation over the nine years (logarithmic scale). (b) Unserved critical load and water shortage for each year and each strategy (logarithmic scales).

With perfect forecasts, S4 is at 0.5 ± 0.3% of the oracle. The 8.5% gap therefore originates almost entirely from forecast errors, and not from the 48 h horizon or from the receding-horizon resolution.

The two other rule-based strategies behave differently. The rule-based MPC (R) protects the critical load (0.24 MWh/year) but leaves 11.5% of water shortage, because pumping is forbidden during any forecast deficit (night and cloudy days) and discharge is forbidden during a forecast surplus. It is worse than S4 in all nine years and worse than S2. The default rule (S1) has 33% water shortage, since the battery never feeds the station. Fixing modes or priorities without optimising the use of the battery for water supply costs more than the served critical energy suggests.

3.2 Trajectories

Figure 4 shows five days poor in renewable energy (February 2016). After a sunny day, S2 fills the tank to 100% with the surplus. When the tank falls below its threshold, S2 feeds the station from the battery. The state of charge thus reaches 15% in the middle of the window, and critical load is not served (466 kWh over the window, against 41 kWh for S4 and 0 for the oracle). S4 delays pumping, keeps the battery for the night and accepts more water

shortage (822 kWh against 493 kWh for S2; 842 kWh for the oracle). The oracle's state-of-charge trajectory is lower than that of S4 at the start of the window. Without a holding cost, the linear program is indifferent to the timing of charge, and the oracle's mean state of charge (37%) is not a realistic target; it explains its 18-year lifetime (Section 3.5).

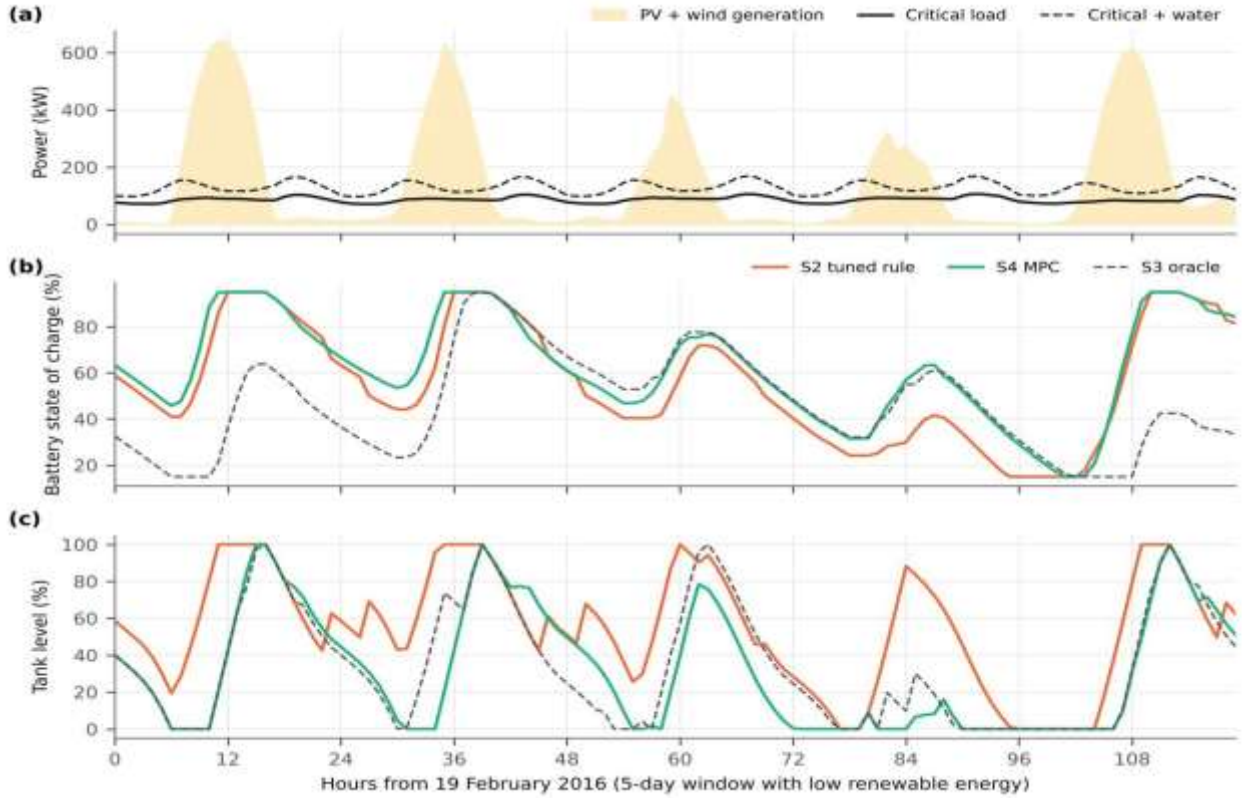


Figure 4. Five days of February 2016: (a) production and loads; (b) battery state of charge; (c) tank level, for the tuned rule S2, the predictive control S4 and the oracle S3.

3.3 Forecast errors, horizon and terminal values

The sensitivities of Sections 3.3 to 3.6 cover three years (2016, 2020, 2024), and S2 remains tuned on 2015. Over these three years, the excess of J of the predictive controller S4 over the oracle is 8.1 ± 1.7%, against 56.7 ± 1.6% for S2 and 0.8 ± 0.3% for S4 with perfect forecasts. Figure 5a multiplies the amplitude of the forecast errors (PV, wind, critical load) by a factor from 0.5 to 4. The excess grows from 4.5% to 40.7 ± 8.1%, about 10 points per unit of scale, while unserved critical load goes from 0.29 to 1.13 MWh/year and levels off beyond a factor of 3. Within the tested range, predictive control remains better than the tuned rule S2. We did not look for the factor at which the two cross, which lies beyond 4.

Figure 5b varies the horizon H. With 6 h, the excess reaches 124.8 ± 13.7% and unserved critical load 10.6 MWh/year, almost at the level of the strategy without flexibility (12.6 MWh/year). Water shortage is the lowest there (0.5%), which we read as a priority given to filling the tank at the expense of the energy stored for the critical load. At 12 h the excess (53.2%) is comparable to that of S2, and it falls to 16.5% at 24 h, 8.1% at 48 h and 6.5% at 96 h. A horizon of at least 24 h is therefore necessary, and the gain beyond 48 h is small.

Persistence forecast. To place the simulated errors against a simple operational forecast, S4 was rerun over the nine years with the 24 h persistence forecast described in Section 2.5 instead of the simulated errors. The excess over the oracle is then 59.4 ± 10.3% (95% interval over years: 53.2–65.7%), compared with 8.5 ± 1.4% with the nominal simulated errors and 53.4 ± 3.1% for the tuned rule S2. S4 with persistence has a higher J than S2 in five of the nine years (4% higher on average) and a lower J than the strategy without flexibility in all nine. Unserved critical load stays low (0.60 ± 0.50 MWh/year, against 0.54 for S4 with simulated errors and 3.33 for S2), but water shortage rises to 4.4% (2.3% for S4 and 2.1% for S2). With a naive forecast, predictive control therefore keeps its protection of the critical load but loses its advantage on criterion J. The benefit reported above depends on forecast quality, and the nominal simulated errors should be read as the accuracy a good forecasting system would have to reach, not as what persistence provides.

The terminal value granted to stored energy and to tank water at the end of the window has a limited effect: the excess is 7.3% with (vb, vw) = (0, 0), 8.1% with the nominal values (0.15; 0.30) and 8.1% when doubling them.

Over a 48 h horizon the solution depends little on this setting, which limits the risk of an arbitrary parameterisation.

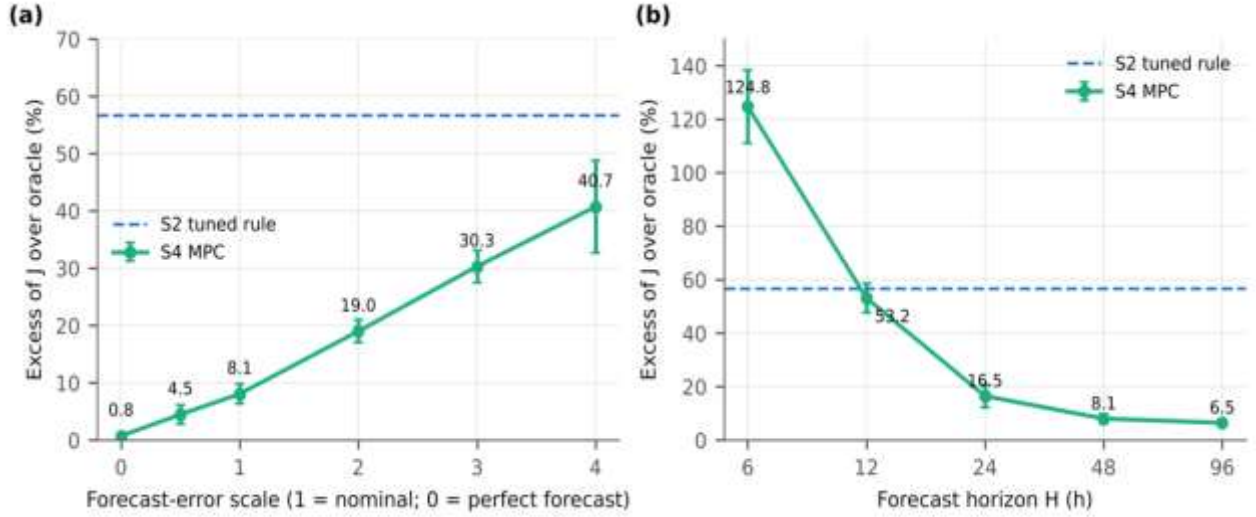


Figure 5. Excess of criterion J over the oracle (mean ± standard deviation, 2016, 2020 and 2024): (a) as a function of the forecast-error scale (0: perfect forecast); (b) as a function of the forecast horizon. The dashed line gives S2.

3.4 Value of flexibility: tank size and deferrable share

Figure 6a varies the tank capacity from 3 to 48 h of mean demand. Each strategy is compared with the oracle of the same configuration, and S2 is again tuned on 2015 for each configuration. Without flexibility (NF), the excess over the oracle increases with capacity (112% at 3 h, 466% at 48 h), because the oracle takes advantage of the tank and NF does not. The tuned rule S2 goes from 92% to 31%. Predictive control stays between 8 and 17%: 14.2% at 3 h, 11.8% at 6 h, 8.1% at 12 h, 9.3% at 24 h and 16.8% at 48 h, with a growing interannual dispersion for large tanks (standard deviations of 5.2 and 6.0 points at 24 h and 48 h). We do not try to explain the rise at 48 h, which is not monotonic and rests on three years.

Figure 6b varies the share of the water load in total energy (15, 35 and 50%). Predictive control stays around 10% (9.8; 8.1; 9.9%), S2 between 37 and 57% and NF between 126 and 272%. With a zero share, the whole load is critical, the three strategies are identical and the excess is zero; this case contains no flexibility to exploit and is not plotted. In that case the chosen sizing is no longer suitable (34.8 MWh/year of unserved critical load), so the value of prediction depends on both the available flexibility and the sizing.

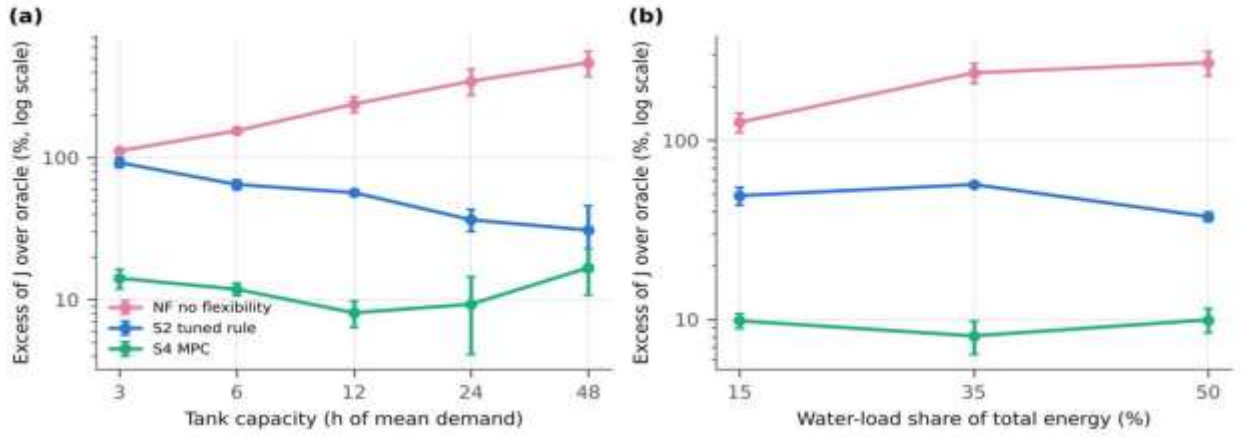


Figure 6. Excess of criterion J over the oracle of the same configuration (logarithmic scale; mean ± standard deviation over 2016, 2020 and 2024): (a) tank capacity; (b) share of the water load.

3.5 Battery ageing and state-of-charge holding cost

Estimated lifetime is 10.8 ± 0.1 years for predictive control against 11.6 ± 0.2 years for S2 (nine years). Figure 7b breaks down the annual damage over the three sensitivity years. It is dominated by calendar ageing (about 0.089 per year for S4 against 0.055 for the oracle), and the share due to cycling (rainflow) is of the order of 0.003 per year. The lifetime gap therefore comes from the mean state of charge, which is higher with predictive control (0.69 against 0.66 for S2; 0.38 for the oracle). The oracle's lifetime (17.9 years over the three sensitivity years, 18.0 years

over the nine years of Table 2) comes from its low mean state of charge (0.38). This low value is an effect of the degeneracy of the linear problem's solution (several trajectories of equal cost) and not a realistic target; we do not present it as an attainable gain.

To test a trade-off, strategy S5 adds a holding cost ch·Eb on the stored energy to J. Over nine years, with ch = $3 \cdot 10^{-3}$ p.u./kWh, the excess of J becomes 27.8 ± 4.0%, unserved critical load 1.51 MWh/year (0.54 for S4, 3.33 for S2) and lifetime 11.6 ± 0.2 years. Figure 7a shows the three sensitivity years: from ch = 10^{-4} to 10^{-2}, lifetime goes from 11.2 to 11.7 years, i.e. to the level of S2, while unserved critical load goes from 1.33 to 2.11 MWh/year and the excess of J from 13.4% to 33.8%. A sufficient holding cost therefore brings lifetime back to that of S2, but at the expense of reliability and of criterion J. At this level it still has better reliability than S2 over the same three years (2.11 against 3.87 MWh/year; the nine-year value for S2 is 3.33 MWh/year), with no lifetime gain. In this model, where damage is essentially calendar-based and temperature is not modelled, the ageing lever is weak: the holding cost is a setting, not a free gain.

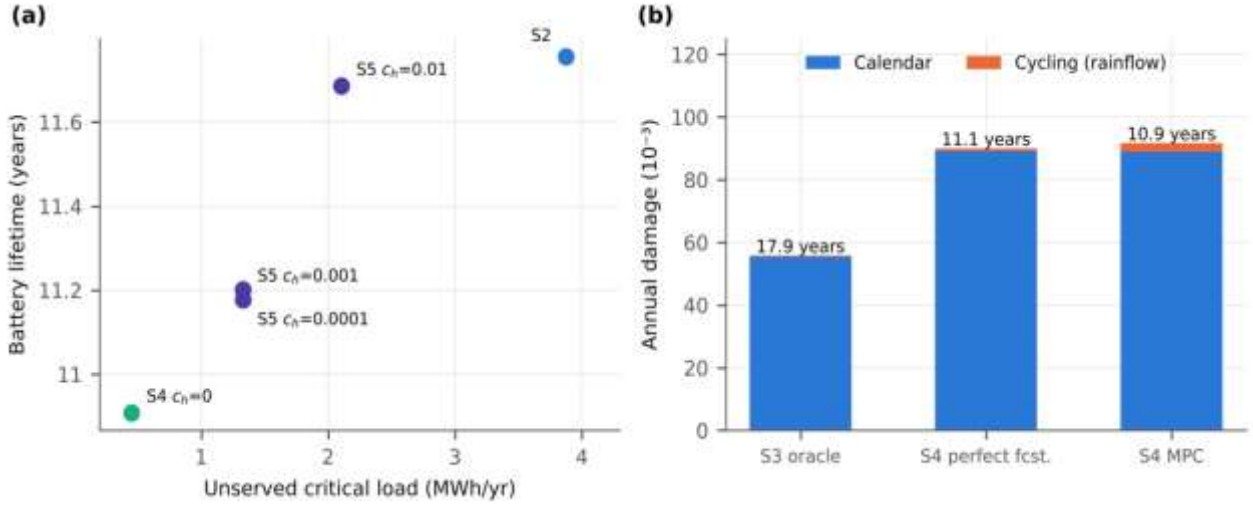


Figure 7. Ageing (2016, 2020 and 2024): (a) lifetime as a function of unserved critical load for S2, S4 and S5 (ch in p.u./kWh); (b) annual calendar and cycling damage for the oracle S3, S4 with perfect forecasts and S4 (flat MPC).

3.6 Robustness to the choice of weights

The weights cun = 3 and cw = 1 p.u./kWh have no monetary meaning. We redid the oracle, NF, S2 (re-tuned on 2015) and S4 with cw = 0.3 and cw = 3 (three years). The excess over the oracle is 363.9 ± 37.0% (NF), 70.4 ± 16.4% (S2) and 20.3 ± 11.0% (S4) for cw = 0.3, and 151.8 ± 60.5%, 20.3 ± 6.1% and 4.6 ± 1.9% for cw = 3. The order NF > S2 > S4 is preserved. For cw = 3 (equal to cun), S4 nevertheless accepts 8.3 MWh/year of unserved critical load against 4.9 for S2, since the criterion then puts water and critical load on the same footing. The result depends on this choice, which should be set by the operator.

3.7 Discussion and limitations

What the results show. In a system without a backup generator, a greedy strategy coincides with the oracle as long as the whole load is critical, so prediction has value only if a deferrable load allows energy to be shifted. With this flexibility, MPC reduces the criterion by 29% relative to the tuned rule and by 68% relative to the strategy without flexibility, and it recovers 84% of the gap between the rule and the oracle. The difference with S2 has the same sign in all nine years. MPC also supplies 21 times less unserved critical load than the strategy without flexibility. The decomposition quantifies the role of prediction: 0.5% excess with perfect forecasts, 8.5% with the nominal errors, and a 24 h horizon as a threshold below which the advantage disappears. These figures hold for the simulated error process. With a 24 h persistence forecast the advantage over the tuned rule on J disappears (Section 3.3), so the value of prediction depends on the accuracy actually attained.

Comparison with a sign rule. Strategy R, inspired by Pippia et al. (2020), uses only the sign of the forecast balance to forbid some actions. It is worse than S4 over the nine years (230 ± 67% excess) because it forbids discharge during forecast surpluses and thus blocks actions that are useful for the tank. Prediction is useful when it serves to price the trade-off, not only to restrict it. This result holds for our implementation and our system, not for the original method.

Relation to the literature. The direction of our forecast result agrees with the one reported by Houben et al. (2023): predictive dispatch loses much of its advantage when perfect foresight is replaced by real forecasts. The two settings differ (a grid-connected community judged on cost, against an islanded site judged on a dimensionless criterion), so the numbers are not comparable. Both nonetheless point to forecast quality as the main limit once the controller is well formulated. The gain over non-predictive operation that Manzano et al. (2022) obtained with

real data is also consistent with the 29% reduction over a tuned rule found here, again with a different system and metric. On ageing, our finding that calendar damage dominates is a property of our assumed damage model. Mazidi et al. (2023) and Sowe et al. (2022) support the broader point that scheduling and current limits should take battery wear into account, which our holding cost S5 does only in a crude way.

Reliability and ageing. The gain of MPC comes mainly from the critical load: 0.54 MWh/year unserved against 3.33 for S2. It is paid for by a higher state of charge and therefore, in our model where ageing is essentially calendar-based (Downing & Socie, 1982; Keil et al., 2016), by a lifetime shorter by 0.8 years. This effect depends on the calendar damage model and on the absence of a thermal model, a limitation that is particularly relevant in the Sahara. It should be seen as a working hypothesis.

Limitations. (i) The load is a published national profile (Load Consumption Data Algeria, 2023), scaled, and not a site measurement; the split into critical and water load, as well as the water profile (two peaks at 7 h and 19 h, seasonal variation of ±12.5%), are assumptions. (ii) Forecast errors are simulated by a multiplicative AR(1) process ($\rho = 0.9$), with one draw per year in the main results (five additional draws change the mean excess by less than one point, Section 3.1), apart from a persistence reference (Section 3.3). The simulated errors do not come from a validated weather-forecast model, and the water load is assumed known. (iii) The pump is continuous, without minimum power or start-up cost, and the tank is lossless. (iv) The battery model ignores temperature, sizing is not optimised, and the PV chain is not validated against measured output. A measured plant in the Algerian Sahara shows an annual performance ratio of 67% (Chaib et al., 2024), lower than the 0.92 system efficiency used here, although the two quantities are not strictly the same. (v) The criterion weights are in p.u./kWh, without monetary value, and the result depends on them in part (Section 3.6); no cost study is made. (vi) The sensitivities of Sections 3.3 to 3.6 rest on three years. (vii) There is no experimental or hardware-in-the-loop validation. (viii) The diesel generator and the hydrogen chain are excluded; they will be the subject of a second article.

4. Conclusion

On a Saharan PV–wind–battery microgrid with a deferrable water load, linear predictive control with a 48 h horizon stays at 8.5 ± 1.4% of the oracle criterion, against 53.4 ± 3.1% for the best rule we tuned and 242 ± 33% without flexibility. Forecast error explains almost all of the remaining gap (0.5% with perfect forecasts). A horizon of at least 24 h is required in this configuration, and the advantage persists up to four times the nominal simulated error. With a naive 24 h persistence forecast, however, the excess rises to 59.4 ± 10.3%, no better than the tuned rule, although the critical load remains well protected. Terminal values have a limited effect. The controller keeps the battery more charged, with an estimated lifetime shorter by 0.8 years, and a state-of-charge holding cost brings this duration back to the level of the rule at the expense of reliability. These results hold for the load, component models and forecast-error assumptions stated above. Extending the assessment to a site-measured load, forecasts from a validated operational model and hardware-in-the-loop tests would allow the reported gaps to be confirmed under field conditions.

References

1. Amrit, R., Rawlings, J. B., & Angeli, D. (2011). Economic optimization using model predictive control with a terminal cost. Annual Reviews in Control, 35(2), 178–186. https://doi.org/10.1016/j.arcontrol.2011.10.011
2. Boutaghane, A., Aksas, M., Chabane, D., & Lebaal, N. (2025). Feasibility and sensitivity analysis of an off-grid PV/wind hybrid energy system integrated with green hydrogen production: A case study of Algeria. Hydrogen, 6(4), 103. https://doi.org/10.3390/hydrogen6040103
3. Cecilia, A., Carroquino, J., Roda, V., Costa-Castelló, R., & Barreras, F. (2020). Optimal energy management in a standalone microgrid, with photovoltaic generation, short-term storage, and hydrogen production. Energies, 13(6), 1454. https://doi.org/10.3390/en13061454
4. Chaib, M., Dahbi, A., & Benatiallah, A. (2024). Evaluation of the operational efficiency of 3 MWp photovoltaic power plant situated in the Sahara Desert of Algeria. International Journal of Power Electronics and Drive Systems, 15(4), 2618–2630. https://doi.org/10.11591/ijpeds.v15.i4.pp2618-2630
5. Downing, S. D., & Socie, D. F. (1982). Simple rainflow counting algorithms. International Journal of Fatigue, 4(1), 31–40.
6. Elsir, M., Al-Awami, A. T., Antar, M. A., Oikonomou, K., & Parvania, M. (2023). Risk-based operation coordination of water desalination and renewable-rich power systems. IEEE Transactions on Power Systems, 38(2), 1162–1175. https://doi.org/10.1109/TPWRS.2022.3174565
7. Gholami, M., & Pisano, A. (2022). Model predictive operation control of islanded microgrids under nonlinear conversion losses of storage units. Electricity, 3(1), 33–50. https://doi.org/10.3390/electricity3010003
8. Houben, N., Cosic, A., Stadler, M., Mansoor, M., Zellinger, M., Auer, H., Ajanovic, A., & Haas, R. (2023). Optimal dispatch of a multi-energy system microgrid under uncertainty: A renewable energy community in Austria. Applied Energy, 337, 120913. https://doi.org/10.1016/j.apenergy.2023.120913

9. Hu, J., Shan, Y., Yang, Y., Parisio, A., Li, Y., Amjady, N., Islam, S., Cheng, K. W., Guerrero, J. M., & Rodriguez, J. (2024). Economic model predictive control for microgrid optimization: A review. IEEE Transactions on Smart Grid, 15(1), 472–484. https://doi.org/10.1109/TSG.2023.3266253
10. Huangfu, Q., & Hall, J. A. J. (2018). Parallelizing the dual revised simplex method. Mathematical Programming Computation, 10(1), 119–142. https://doi.org/10.1007/s12532-017-0130-5
11. Keil, P., Schuster, S. F., Wilhelm, J., Travi, J., Hauser, A., Karl, R. C., & Jossen, A. (2016). Calendar aging of lithium-ion batteries I. Impact of the graphite anode on capacity fade. Journal of The Electrochemical Society, 163(9), A1872–A1880. https://doi.org/10.1149/2.0411609jes
12. Kimber, A., Mitchell, L., Nogradi, S., & Wenger, H. (2006). The effect of soiling on large grid-connected photovoltaic systems in California and the Southwest region of the United States. In 2006 IEEE 4th World Conference on Photovoltaic Energy Conversion. https://doi.org/10.1109/wcpec.2006.279690
13. Load Consumption Data Algeria. (2023). Load consumption data Algeria [Data set]. Mendeley Data. https://doi.org/10.17632/z5x2d3mhw7.1 — see also: Data driven models on load forecasting: Case study Algeria, Data in Brief. https://doi.org/10.1016/j.dib.2023.109854
14. Manzano, J. M., Salvador, J. R., Romaine, J. B., & Alvarado-Barrios, L. (2022). Economic predictive control for isolated microgrids based on real world demand/renewable energy data and forecast errors. Renewable Energy, 194, 647–658. https://doi.org/10.1016/j.renene.2022.05.103
15. Mayne, D. Q., Rawlings, J. B., Rao, C. V., & Scokaert, P. O. M. (2000). Constrained model predictive control: Stability and optimality. Automatica, 36(6), 789–814. https://doi.org/10.1016/S0005-1098(99)00214-9
16. Mazidi, M., Khezri, R., Mohiti Ardakani, M., Le, A. T., & Steen, D. (2023). Effects of calendar and cycle ageing on battery scheduling for optimal energy management: A case study of HSB Living Lab. In 2023 IEEE International Conference on Energy Technologies for Future Grids (ETFG). https://doi.org/10.1109/ETFG55873.2023.10407383
17. Mokhtara, C., Negrou, B., Bouferrouk, A., Yao, Y., Settou, N., & Ramadan, M. (2020). Integrated supply–demand energy management for optimal design of off-grid hybrid renewable energy systems for residential electrification in arid climates. Energy Conversion and Management, 221, 113192. https://doi.org/10.1016/j.enconman.2020.113192
18. Mueller, M. A., Angeli, D., & Allgöwer, F. (2014). On the performance of economic model predictive control with self-tuning terminal cost. Journal of Process Control, 24(8), 1179–1186. https://doi.org/10.1016/j.jprocont.2014.05.009
19. NASA POWER Project. (2026). NASA Prediction of Worldwide Energy Resources (POWER): hourly data for Ouargla (31.95° N, 5.33° E), 2015–2024 [Data set]. NASA Langley Research Center. https://power.larc.nasa.gov/ (accessed 9 October 2026).
20. Pippia, T., Sijs, J., & De Schutter, B. (2020). A single-level rule-based model predictive control approach for energy management of grid-connected microgrids. IEEE Transactions on Control Systems Technology, 28(6), 2364–2376. https://doi.org/10.1109/TCST.2019.2945023
21. Skoplaki, E., & Palyvos, J. A. (2009). On the temperature dependence of photovoltaic module electrical performance: A review of efficiency/power correlations. Solar Energy, 83(5), 614–624. https://doi.org/10.1016/j.solener.2008.10.008
22. Sowe, J., Varela Barreras, J., Schimpe, M., Wu, B., & Candelise, C. (2022). Model-informed battery current derating strategies: Simple methods to extend battery lifetime in islanded mini-grids. Journal of Energy Storage, 51, 104524. https://doi.org/10.1016/j.est.2022.104524
23. Wang, S., Guo, D., Han, X., Lu, L., Sun, K., Li, W., Sauer, D. U., & Ouyang, M. (2020). Impact of battery degradation models on energy management of a grid-connected DC microgrid. Energy, 207, 118228. https://doi.org/10.1016/j.energy.2020.118228
24. Wicke, M., & Bocklisch, T. (2024). Hierarchical energy management of hybrid battery storage systems for PV capacity firming and spot market trading considering degradation costs. IEEE Access, 12, 52669–52686. https://doi.org/10.1109/ACCESS.2024.3387748
25. Zhang, Y., Fu, L., Zhu, W., Bao, X., & Liu, C. (2018). Robust model predictive control for optimal energy management of island microgrids with uncertainties. Energy, 164, 1229–1241. https://doi.org/10.1016/j.energy.2018.08.200