Is a 100% renewable electricity mix cost-efficient under weather uncertainty?

Job Market Paper - Latest version here

Konrad Dierks*

4 November 2026

Abstract

Renewables (wind and solar) do not emit greenhouse gases, but their output is highly uncertain and intermittent. This raises the question of how costly a 100% renewable electricity mix would be, and how this cost depends on the cost and efficiency of the available storage technologies. I study these questions in a dynamic stochastic model of wind and solar power generation and storage. The model is calibrated on German data and accounts for both flexible consumption and demand curtailment. I first solve numerically for the optimal operation of given generation and storage capacities, using a semi-Lagrangian scheme with relative value iteration. I then optimize over renewable and storage capacities to obtain the efficient electricity mix. With battery storage only and a 5% discount rate, the optimal 100% renewable mix requires multiplying the 2025 wind and solar capacities about sevenfold, at a welfare-adjusted levelized cost of 285 €/MWh. Allowing for lossy, yet inexpensive, hydrogen storage reduces the optimal renewable capacities sharply and lowers this cost to 142 €/MWh. Relative to perfect foresight over the observed decade, weather uncertainty mainly raises the optimal hydrogen storage capacity, from 23 to 35 TWh. With a fossil backup, a non-negligible storage capacity becomes optimal only at a positive carbon price, and marginal abatement costs rise sharply close to zero emissions if there are no other low-carbon alternatives.

*PhD student, Toulouse School of Economics, University of Toulouse Capitole

This work was supported by two French government grants managed by the Agence Nationale de la Recherche under the "Investissements d'avenir" program (references "ANR-18-EURE-0021", "ANR-21-ESRE-0051"). Under the latter, this work was granted access to the MesoNET resources center and the MesoNET Project under the allocation m26158. Furthermore, this work was supported by the Agence Nationale de la Recherche through the AVINEQ project, led by Christian Gollier.

1 Introduction

Decarbonizing electricity supply requires replacing fossil generation while maintaining reliable access to electricity. Germany has committed by law to net greenhouse-gas neutrality by 2045 (Deutscher Bundestag, 2019), and the European Union to climate neutrality by 2050 (European Parliament, 2021). In Germany, wind and solar power are central to this transition, but their output varies with the weather and the time of day, and is subject to uncertainty. Storage can shift surplus production to periods of scarcity.

This paper asks how costly a 100% renewable electricity mix would be when dispatch decisions must be made under weather uncertainty, and how this cost depends on the cost and efficiency of the available storage technologies. To answer these questions, I determine the combination of wind, solar, and storage capacities that maximizes welfare. I study the German electricity system in two settings: a fully renewable mix with battery storage only or with both battery and hydrogen storage, and a mix that also allows fossil backup at increasing social costs of carbon. This comparison distinguishes the cost of a zero-emission system from the marginal cost of eliminating the remaining fossil generation. I also ask how much weather uncertainty changes the optimal mix and its operating value relative to perfect foresight of the weather. The analysis proceeds in two steps. First, I solve the stochastic control problem of operating the electricity system for any given capacity mix. Second, I compare the resulting marginal values of wind, solar, battery, and hydrogen capacity with their marginal costs to characterize the optimal mix.

Solar and wind power have become highly competitive for decarbonizing electricity supply. In Germany, a rising renewable share has cut emissions from electricity production by more than half since 1990 (Icha and Lauf, 2025), despite the nuclear exit completed in 2023. Flexible gas plants can cover renewable shortfalls, yet continued reliance on them entails carbon emissions and exposure to fossil-fuel supply shocks. Doubts about the true cost and feasibility of a fully renewable system stem from the difficulty of managing intermittency. The limited response of electricity demand to prices implies that, absent storage, a fully renewable system built on wind and solar would regularly fail to balance supply and demand, causing costly blackouts. Generation overcapacity, flexible consumption, and electricity storage can each help manage this volatility, but each carries its own welfare cost.

Storage makes this balance an intertemporal problem: electricity consumed today is no longer available to cover a future shortfall. Evaluating storage dispatch on historical weather

paths can obscure parts of this trade-off if the operator is assumed to know future renewable production. The value of preserving stored energy depends on the risk of future scarcity. Wind and solar production can be forecast only a few days or weeks ahead. In electricity markets, the uncertainty of renewable production can matter more than its intermittency (Weber and Woerman, 2024). With a single storage technology, fixed demand, and a perfectly flexible, unconstrained fossil backup, dispatch follows a simple rule: surplus production is stored, storage is drawn down when production falls short, and the backup takes over once storage is empty. With several storage technologies, price-responsive demand, or an absent, constrained, or insufficiently flexible backup, optimal dispatch requires a model of the uncertainty of future renewable production.

Approach. Capacities are chosen before weather is realized, accounting for optimal operation of the installed system under uncertainty. All capacities are chosen at their full lifecycle cost, so existing assets do not enter the mix, and Germany is modeled without electricity trade with neighboring countries and without transmission or ramping constraints. For each capacity mix, I solve a continuous-time stochastic control problem using a calibrated joint stochastic process for German wind and solar production, including a persistent, seasonal drought regime, together with predictable intraday and seasonal variation in production and demand. The state of the system consists of the battery and hydrogen storage levels; the wind and solar weather state and the drought regime, which are the only sources of production forecasts; and the hour of the year. Optimal dispatch responds to current weather conditions and storage levels while respecting storage energy and power limits. It maximizes the expected long-run average benefit of electricity consumption net of fuel costs; since this benefit is concave in consumption, volatile consumption lowers welfare. The expected operating value under optimal dispatch, net of lifecycle capacity costs, determines the welfare value of the mix. Mixes are compared through a welfare-adjusted levelized cost: system cost plus the welfare loss from consumption below its reference level, net of the welfare gain above it, per MWh of reference consumption. Minimizing it is equivalent to maximizing welfare.

The two storage technologies involve different trade-offs. Batteries have a high round-trip efficiency of 0.91 but expensive energy capacity, at an investment cost of about 350 €/kWh, making prolonged renewable shortfalls costly to cover. Hydrogen storage in salt caverns ("power-to-gas-to-power") offers much cheaper energy capacity, at about 3.7 €/kWh, but converting electricity to hydrogen and back requires costly equipment and entails a round-trip efficiency of only 0.34.

Adapting consumption to availability is a key ingredient of an efficient electricity system with a high renewable share. Welfare accounts for flexible consumption, calibrated to the short-run price elasticity of German electricity demand, and for the cost of demand curtailment once price-sensitive consumers cannot adjust further.

Main findings. First, batteries alone cannot solve the intermittency problem of a fully renewable mix, even with the calibrated consumption flexibility: the quantities required for full decarbonization without major supply outages are too costly. Restricting storage to batteries leads to more volatile consumption under optimal dispatch and a higher welfare-adjusted levelized cost. At a 5% discount rate, adding hydrogen to a battery-only mix reduces the welfare-adjusted levelized cost by about half. Even with low battery costs, the optimal battery-only mix has a welfare-adjusted levelized cost of 211 €/MWh, against 136 €/MWh when hydrogen storage is also available. Conversely, removing batteries from the mix with hydrogen storage raises the welfare-adjusted levelized cost by only about 1%: cheap long-duration storage is crucial for a cost-efficient fully renewable mix.

Second, adding hydrogen reduces the renewable capacity required: optimal wind capacity falls from 547 GW to 312 GW and solar capacity from 601 GW to 216 GW. With batteries alone, additional renewable capacity is needed to sustain production during low-wind, low-solar periods, known in German as *Dunkelflaute* (Kittel and Schill, 2024), even though it produces large surpluses at other times. This additional capacity raises production in low-production periods more cheaply than a large battery capacity. Cheap hydrogen energy capacity reduces the need for this additional generation despite its lower round-trip efficiency. Renewable overcapacity and storage are therefore substitutes in this setting. "Wasting" renewable production is more common in the battery-only mix, and avoiding it is not welfare-optimal.[1]

Third, the optimal mix with both storage technologies at a 5% discount rate contains around 35 TWh of hydrogen storage and 44 GWh of battery storage, compared to around 60 GWh of installed pumped storage and battery storage combined in 2025. Its welfare-adjusted levelized cost lies within the range, from 121 to 152 €/MWh, of the welfare-adjusted levelized costs of meeting the same demand with natural gas, hard coal, or lignite alone at recent fuel prices and the August 2026 carbon price (Appendix O). It exceeds that of the optimal mix with fossil backup by about 17 €/MWh at a carbon price of 750 €/tCO_2 and by about 50 €/MWh without a carbon price (Section 7). At a 2% discount rate, it falls substantially, from 142 to 97 €/MWh, which highlights the importance of the cost of capital.

[1] This contrasts with the zero-waste constraint in some of the storage scenarios of Sinn (2017).

Fourth, the optimal fully renewable mix does not fully smooth consumption: in states with low renewable production and depleted storage, consumption falls below its reference level, which implies a welfare cost. The optimal mix trades off the capital cost of the storage system against this welfare cost. Greater demand flexibility has a moderate effect on welfare and the welfare-adjusted levelized cost. It lowers optimal wind capacity and substantially reduces hydrogen storage, while solar and battery capacities change little: more flexible demand reduces the need for insurance against low-production events.

Fifth, allowing fossil backup reveals a distinction between average electricity costs and marginal abatement costs. Without carbon pricing, renewables are cost-efficient even with virtually no storage, since fossil plants cover low-wind, low-solar states. A positive carbon price is needed for further decarbonization because storage is costly relative to the fossil alternative. Hydrogen storage enters the optimal mix as the carbon price rises and further emissions reductions become more costly. At a carbon price of about 300 €/tCO_2, the optimal mix cuts emissions by about 95% relative to 1990, with 2.6 TWh of hydrogen storage, at a welfare-adjusted levelized cost of 116 €/MWh. Near zero emissions, however, the marginal abatement cost exceeds 750 €/tCO_2 when no carbon-neutral backup fuel (e.g., biofuels or imported hydrogen) is available. Rare fossil generation remains valuable during scarcity.

Finally, comparing the operating value under uncertainty with that under perfect foresight reveals that perfect foresight of renewable production has low value compared to the cost of the optimal mix, because the optimal mix already holds a large hydrogen stockpile, which cheap cavern storage capacity makes affordable. Moreover, the optimal mix under perfect foresight over a 10-year weather realization has renewable and conversion capacities similar to those of the optimal mix under weather uncertainty, but substantially less hydrogen storage. Knowing the weather in advance therefore mainly reduces the stockpile held as insurance against low-production periods: for 2016-2025, the perfect-foresight optimum holds 23 rather than 35 TWh of hydrogen storage, perfect foresight is worth more than twice as much to it, and under weather uncertainty it performs worse than the optimal mix under uncertainty by about 2.7 billion euros per year (B€/y). However, the variation in the perfect-foresight optimal mix is substantial when its underlying weather realizations consist of single observed years. The optimal storage operation can differ substantially on the same weather realization under perfect foresight and weather uncertainty.

Literature and contributions. The paper contributes to the literature on renewable intermittency and decarbonization (Gowrisankaran, Reynolds and Samano, 2016; Sinn, 2017; Helm

and Mier, 2019; Ambec and Crampes, 2019; Holland, Mansur and Yates, 2022) by making, to my knowledge, for the first time, wind, solar, battery, and hydrogen capacities jointly endogenous while dispatch is optimal under weather uncertainty at every instant and consumption responds to scarcity. The closest papers relax one or more of these elements: Brown, Neumann and Riepin (2025) combine them with elastic demand but choose capacities under perfect foresight; Schmidt (2025) and Blanchard and Megy (2026) choose a similar set under uncertainty but with inelastic demand and perfect foresight within monthly stages.

Geske and Green (2020), Shen, Ilic and Parsons (2026), and Keske, Steffen and Parsons (2026) link precautionary storage operation under uncertain residual demand, net demand, or renewable inflow to investment in conventional or thermal capacity. They take renewable capacity as given and, except for Shen, Ilic and Parsons (2026), storage capacity as well. I make renewable and storage capacities jointly endogenous, so that precautionary operation shapes the renewable-storage mix itself.

The strand most directly related to this paper sizes long-duration storage for reliability in high-renewable systems, with a focus on droughts, mostly with perfect foresight of historical weather (Ziegler et al., 2019; Dowling et al., 2020; Sepulveda et al., 2021; Ruhnau and Qvist, 2022; Kittel, Roth and Schill, 2026). Ruhnau and Qvist (2022) jointly choose wind, solar, battery, and hydrogen capacities with fixed demand and perfect foresight over 35 weather years. My optimal mix has a qualitatively similar composition, with a small battery energy capacity beside a much larger hydrogen storage capacity.[2] Relative to perfect-foresight sizing studies, I quantify the precautionary premium for long-duration storage: weather uncertainty raises the optimal hydrogen stockpile from 23 to 35 TWh, and a mix sized with perfect foresight loses about 2.7 B€/y when operated under uncertainty (Section 6). This extends the precautionary-storage mechanism of Geske and Green (2020) to long-duration hydrogen storage.

The paper also relates to work on storage valuation and capacity planning (Steffen and Weber, 2013; Sisternes, Jenkins and Botterud, 2016; Zerrahn and Schill, 2017; Zerrahn, Schill and Kemfert, 2018; Crampes and Trochet, 2019; Junge, Mallapragada and Schmalensee, 2022), and to work on market structure, policy, and the energy transition (Helm and Mier, 2021; Pommeret and Schubert, 2022; Andrés-Cerezo and Fabra, 2023; Gentile, 2024; Antweiler and Muesgens, 2025). Zerrahn, Schill and Kemfert (2018) show that allowing cost-optimal curtail-

[2]The optimal hydrogen storage capacity, at 35 TWh, turns out to be smaller than their 55 TWh even under weather uncertainty, likely because Ruhnau and Qvist (2022) assume inflexible demand and lower capacity costs for hydrogen storage, and because the weather model may not reproduce the full variation of renewable production across years (Appendix E).

ment of renewable surpluses, rather than imposing a zero-waste constraint, greatly reduces storage requirements. Antweiler and Muesgens (2025) show that generation and storage can be substitutes or complements depending on the technology pair and the correlation between supply and demand, Pommeret and Schubert (2022) find that renewable generation and storage are complements, and Blanchard and Megy (2026) find that battery and hydrogen storage are imperfect substitutes. I find that the substitution between renewable overcapacity and storage persists under uncertainty as insurance against renewable droughts, and that the battery-solar relation flips from substitute to complement once hydrogen storage is available.[3] The marginal values of the two storage technologies are strongly asymmetric: adding hydrogen storage to a battery-only mix halves the welfare-adjusted levelized cost, whereas removing batteries from the mix with hydrogen storage raises it by only about 1%.

Many capacity models take demand as fixed. Brown, Neumann and Riepin (2025) combine wind, solar, battery, and hydrogen capacities with elastic demand in a long-term capacity expansion under perfect foresight, and use a 96-hour rolling-horizon operational model with fixed capacities to study storage bidding and price formation. I instead choose capacities under weather uncertainty with scarcity-responsive consumption, so that the optimal mix itself reflects the welfare cost of uninsured weather risk.

I also contribute to the literature on the cost of deep decarbonization (Sepulveda et al., 2018; Holland, Mansur and Yates, 2022) by quantifying the cost of the last emission reductions in a stochastic model with endogenous battery and hydrogen storage and fossil backup: without a carbon-neutral firm resource (one that can supply electricity on demand for as long as needed, independently of the weather), marginal abatement costs exceed 750 €/tCO_2 near zero emissions (Section 7). This extends the firm-resource argument of Sepulveda et al. (2018) to a setting with explicit weather risk and joint storage choice.

Levelized costs of individual technologies ignore that intermittent output is worth less than dispatchable output (Joskow, 2011; Hirth, 2013). I therefore compare complete capacity mixes that serve the same demand, using a welfare-adjusted levelized cost that adds to system cost the net welfare loss from consumption deviating from its reference level. This places the fully renewable mix, fossil-only supply, and mixes with fossil backup on a common welfare basis.

Finally, the model connects to buffer-stock saving and precautionary storage theory (Deaton, 1991; Durmaz, 2016) and to competitive storage models (Scheinkman and Schechtman, 1983), to continuous-time HJB and PDE models of storage (Chen and Forsyth, 2008; Howell et al., 2009;

[3] Adding hydrogen storage to a battery-only mix reduces optimal wind capacity from 547 to 312 GW and solar capacity from 601 to 216 GW.

Carmona and Ludkovski, 2010; Johnson, Howell and Duck, 2017; Picarelli and Vargiolu, 2021; Bermúdez and Padín, 2024), and to discrete-time dynamic-programming models of electricity storage (Secomandi, 2010; Lai, Margot and Secomandi, 2010; Löhndorf, Wozabal and Minner, 2013; Xi, Sioshansi and Marano, 2014; Asamov, Salas and Powell, 2016; Salas and Powell, 2018). The closest electricity application is Bobtcheff (2011), who combines stochastic renewable supply, consumption utility, constrained hydroelectric storage, and thermal backup. Building on state-constrained stochastic-control approaches such as Picarelli and Vargiolu (2021) and Bermúdez and Padín (2024), I use a Hamilton-Jacobi-Bellman (HJB) equation to characterize dispatch in a calibrated model with wind, solar, batteries, and hydrogen.

Roadmap. Section 2 presents the model. Section 3 describes the data and calibration. Section 4 presents results with batteries as the only storage technology. Section 5 reports the results with both battery and hydrogen storage. Section 6 examines the role of weather uncertainty. Section 7 reports the results with fossil backup and the resulting abatement cost curves. Section 8 discusses the results and their limitations. The appendices provide theoretical and numerical details and additional results.

2 Model

The model has a period-0 investment decision and a period-1 operating problem in continuous time. In period 0, renewable generation capacities, storage energy and power capacities, and, where allowed, fossil backup capacity are chosen to maximize welfare. In period 1, after these capacities are fixed, optimal storage and backup dispatch responds to observed weather conditions, storage levels, and predictable daily and seasonal variation in renewable production and demand. The welfare value of a capacity vector is the present value of the expected long-run average net operating benefit under the optimal period-1 policy, net of period-0 lifecycle capacity costs. This criterion looks at the operating benefit of the system in its stationary state: it abstracts from the transition from initial storage levels and weather conditions and from the timing of benefits within the year.

2.1 Capacities and Weather

A capacity vector q contains renewable capacities, storage energy capacities, storage power capacities, and fossil backup capacity.

$$q = \big(K_W, K_S, \bar{S}_H, \bar{S}_B, K_{H-}, K_{H+}, K_{B-}, K_{B+}, F\big), \tag{1}$$

where K_W and K_S are nominal wind and solar capacities in GW, $\bar{S}_H$ and $\bar{S}_B$ are hydrogen and battery energy capacities in GWh, K_{H-} and K_{B-} are discharging power capacities in GW, K_{H+} and K_{B+} are charging power capacities in GW, and F is fossil backup capacity in GW. The fully renewable case is the restriction $F = 0$. For fixed capacities q, the operating problem is an infinite-horizon stochastic control problem in continuous time. Time is measured in hours. The periodic phase $\theta_t \in [0, T_Y)$ records the time of year, with period $T_Y = 365 \cdot 24 = 8\,760$ hours, and advances at unit speed. Throughout, $\delta > 0$ is the annual discount rate, so a horizon of t hours is discounted by $e^{-\delta t / T_Y}$.

The weather determines the realized capacity factors of renewables, which take values between 0 and 1. As weather fluctuates around stable seasonal patterns, it is natural for weather uncertainty to be driven by a stationary process. I choose a latent two-dimensional Ornstein-Uhlenbeck process, $X_t = (X_{Wt}, X_{St})$, whose mean depends on the drought regime $R_t \in \{0, 1\}$ described below:

$$dX_t = \Psi_X(X_t + R_t m(\theta_t))dt + Q dB_t, \tag{2}$$

where B_t is a two-dimensional Wiener process and $\Psi_X, Q \in \mathbb{R}^{2\times 2}$. The matrices Ψ_X and Q are such that, in the absence of the drought regime, the stationary distribution of X_t would be

bivariate standard normal with independent components; the dependence between wind and solar production enters through the maps below. In the normal regime, $R_t = 0$, the drift pulls the process toward 0, whereas B_t introduces continuous-time Gaussian shocks that can push the process toward or away from this target.

Persistent periods of low renewable production observed in the data[4] are represented by the drought regime R_t, where $R_t = 1$ indicates the scarcity regime. The regime follows a continuous-time Markov chain, independent of B_t, that enters the scarcity regime at rate γ_{01} and leaves it at rate γ_{10}, in the spirit of Markov-switching models (Hamilton, 1989), which have also been applied to wind time series (Ailliot and Monbet, 2012). In the scarcity regime, the drift pulls the latent state toward a season-dependent drought mean $-m(\theta_t)$, whose components are negative. Wind and solar production therefore decline gradually at the onset of a drought and recover gradually at its end. At phase θ_t, the latent state is mapped into the weather state $N_t = n_{\theta_t}(X_t) = (N_{Wt}, N_{St})$. This construction separates the distribution of production for each technology, set by the maps in (3) below, from the dependence between wind and solar production, which is described by a copula C_ϑ (Nelsen, 2006). In the absence of the drought regime, both components of the weather state are standard normal. The copula is a seasonal mixture of a Gaussian copula, which describes the dependence of a bivariate normal distribution and has a seasonally varying correlation, and a Clayton copula; the map and the copula are given in Appendix E. The Clayton copula concentrates positive dependence in the lower tail, so that low values of both components are more closely linked than high values. The mixture allows the dependence between wind and solar production to differ across seasons. These features matter for the capacity choice in distinct ways: persistence and the drought regime determine how much energy must be carried over several days or weeks, and hence the value of storage energy capacity; lower-tail dependence limits the gain from combining wind and solar in the hours of lowest production; and the seasonal variation of production, with abundant solar output in summer and higher average wind output in winter, determines how well a given combination of wind and solar capacity covers demand in each season, and hence how much energy must be shifted between seasons.

Given wind and solar capacities K_W and K_S, renewable production, in GW, is

$$Y_t = K_W \psi_W(\varsigma_W(\theta_t) N_{Wt}, \theta_t) + K_S \psi_S(\varsigma_S(\theta_t) N_{St}, \theta_t) \phi_S(\theta_t). \tag{3}$$

The maps ψ_W and ψ_S transform the scaled components of the weather state into the hourly wind capacity factor and the hourly solar capacity factor before the intraday solar profile is

[4] The autocorrelation structure of a bivariate Ornstein-Uhlenbeck process is exponential and cannot reproduce some of the observed persistence.

applied; ς_W and ς_S are scales that vary over the year and, for wind, over the day, while ϕ_S is the intraday solar profile, which is zero at night. Their calibration is described in Section 3, with functional forms in Appendix E.

2.2 Storage, Dispatch, and Consumption Benefit

There are two storage technologies, indexed by $i \in \{H, B\}$ for hydrogen and batteries[5]. The storage action a_{it}, measured in GW, is positive when electricity is stored and negative when storage is discharged. With $a_{it+} = \max\{a_{it}, 0\}$ and $a_{it-} = \min\{a_{it}, 0\}$,

$$dS_{it} = (\omega_{i+} a_{it+} + \omega_{i-} a_{it-})dt, \qquad 0 \le S_{it} \le \bar{S}_i, \tag{4}$$

where S_{it} and $\bar{S}_i$ are measured in GWh. Write $S_t = (S_{Ht}, S_{Bt})$ for the joint storage-level process. For the corresponding storage component of a state, write $s = (s_H, s_B) \in \left[0, \bar{S}_H\right] \times \left[0, \bar{S}_B\right]$. Power capacities impose $-K_{i-} \le a_{it} \le K_{i+}$. The charging efficiency parameter satisfies $0 < \omega_{i+} \le 1$ and gives the share of forgone consumption that enters storage. The discharging efficiency parameter satisfies $\omega_{i-} \ge 1$ and gives the decrease in stored energy needed to deliver one unit of consumption. Fossil dispatch f_t, measured in GW, respects the fossil backup capacity,

$$0 \le f_t \le F, \tag{5}$$

and has a marginal cost p_f.

Then, the energy-balance constraint is

$$c_t + \sum_{i \in \{H,B\}} a_{it} = Y_t + f_t. \tag{6}$$

Here, c_t is electricity consumption and admissible actions must keep it nonnegative. There is no upper bound on consumption; excess electricity therefore remains in the balance and yields positive but eventually near-zero marginal value.[6] Furthermore, transmission is unconstrained and there are no ramping constraints (i.e., constraints on how fast fossil dispatch or storage actions can change).

Electricity consumption c at calendar phase ϑ generates a phase-dependent benefit $U_\vartheta(c)$, measured in monetary terms. Its derivative $U'_\vartheta(c)$ gives the marginal value of electricity. The consumption-benefit function is specified in Appendix A. Its phase dependence captures variation in demand over the day, the week, and the year through reference consumption $c_{\text{ref}}(\vartheta)$, derived from the observed hourly German load in 2025; see Section 3. Reference consumption

[5] Hydrogen and batteries are the storage technologies presented in the numerical results and modeled through their cost and efficiency parameters. The model can accommodate any set of two storage technologies.

[6] In a real-world electricity grid, renewable generation may instead be curtailed when demand does not respond sufficiently.

can be thought of as the quantity demanded at the phase-independent reference price p_{ref}; the benefit at reference consumption is normalized to zero.

The consumption-benefit function has two branches, which meet at the curtailment threshold $k(\vartheta) = \kappa c_{\text{ref}}(\vartheta)$, a fixed fraction κ of reference consumption. Above the threshold, demand responds to prices: the marginal value of electricity follows the calibrated inverse demand and rises as consumption falls, reaching a marginal value p_{cap} at the threshold. Below the threshold, demand no longer responds to prices, and any further shortfall must be met by curtailing demand. On this curtailment branch, the marginal value is the marginal cost of unserved demand. It rises linearly from p_{cap} at the threshold to $v_{\text{blackout}} > p_{\text{cap}}$ at a full blackout, because partial outages, which affect a small and targeted set of consumers, are less costly per MWh than a full blackout. The benefit and the marginal value are both continuous at the threshold. Because marginal value rises as consumption falls, electricity is worth far more in a shortage than in a surplus. Storage therefore has a precautionary role: the planner accepts conversion losses, idle capacity, and a lower average value of consumption to make deep shortfalls less likely. The price elasticity of demand determines how cheaply moderate shortfalls are absorbed, and the price cap and the blackout value determine the cost of rare deep ones. The consumption-benefit function is calibrated and illustrated in Section 3.

2.3 Operating Problem

An operating state is $\xi = (x, r, s, \vartheta)$, where x is the latent state, r is the drought regime, s contains the storage levels, and ϑ is the calendar phase; write $\Xi_t = (X_t, R_t, S_t, \theta_t)$ for its evolution. The operator observes the full operating state, including the drought regime.[7] The fixed-capacity operating problem maximizes the long-run average hourly net operating benefit:

$$\rho^q = \sup_{(a,f)\in\mathcal{A}} \liminf_{T\to\infty} \frac{1}{T}\mathbb{E}_\xi\left[\int_0^T \left(U_{\theta_t}(c_t) - p_f f_t\right) dt\right], \tag{7}$$

subject to the state dynamics and the storage, dispatch, and energy-balance constraints introduced above. Here $\mathcal{A}$ is the set of admissible controls adapted to the information available at each date.

2.4 Dispatch in the Absence of Storage

The special case without storage makes the operating mechanism transparent. At phase ϑ, renewable output y, and unit fossil cost p_f, fossil dispatch $f^q(\vartheta, y)$ solves

[7] Appendix N compares the results with those of an operator who does not observe the regime and assesses it from the current weather state.

$$f^q(\vartheta, y) \in \arg\max_{0 \le \nu \le F} \big\{ U_\vartheta(y + \nu) - p_f \nu \big\}. \tag{8}$$

Let $\hat{c}_f(\vartheta) = (U'_\vartheta)^{-1}(p_f)$ denote the consumption level at which marginal value equals marginal fossil cost. The optimal dispatch is

$$f^q(\vartheta, y) = \min\big\{ F, \max\big\{ 0, \hat{c}_f(\vartheta) - y \big\} \big\}. \tag{9}$$

Thus fossil generation fills a renewable shortfall until marginal consumption value falls to p_f, unless fossil capacity binds. It is unused when the marginal value of consumption supplied by renewables alone is below p_f. The marginal value of consumption thus follows a merit order with price-responsive demand: it is set by renewable output alone when $y > \hat{c}_f(\vartheta)$, equals p_f when fossil generation is marginal, and exceeds p_f only when fossil capacity binds. The scarcity rents in these last states remunerate fossil capacity.

2.5 Solution of the Dispatch Problem with Storage

Let $W^q(\xi)$ be the relative value of an operating state after accounting for the long-run hourly payoff ρ^q; W^q is defined up to an additive constant. Up to this constant, $W^q(\xi)$ is the expected net operating benefit in excess of ρ^q that accumulates from state ξ under an optimal stationary policy ((35) in Appendix B). It measures the advantage of starting from ξ: it is higher, for example, with more stored energy or outside the scarcity regime. For storage actions $\alpha = (\alpha_H, \alpha_B)$, with $\alpha_{i+} = \max\{\alpha_i, 0\}$ and $\alpha_{i-} = \min\{\alpha_i, 0\}$, fossil action ν, renewable output $y(\xi)$, and storage-energy flow $b_i(\alpha_i) = \omega_{i+}\alpha_{i+} + \omega_{i-}\alpha_{i-}$, the HJB equation of this long-run average problem, or ergodic HJB, is

$$\begin{aligned} \rho^q = {} & \partial_\vartheta W^q + (\Psi_X(x + rm(\vartheta))) \cdot \nabla_x W^q + \frac{1}{2}\,\mathrm{tr}(QQ^T D_x^2 W^q) \\ & + \gamma_{r,1-r}(W^q(x, 1-r, s, \vartheta) - W^q(x, r, s, \vartheta)) \\ & + \sup_{(\alpha,\nu)\in A(\xi)} \left\{ U_\vartheta(y(\xi) + \nu - \alpha_H - \alpha_B) - p_f \nu + \sum_{i \in \{H,B\}} b_i(\alpha_i) \partial_{s_i} W^q \right\}. \end{aligned} \tag{10}$$

Here $A(\xi)$, given in (37) of Appendix B, contains the actions that satisfy the constraints above at state ξ; in particular, empty storage cannot be discharged and full storage cannot be charged. $D_x^2 W^q$ is the Hessian of W^q in the latent state. The supremum balances current consumption benefits net of fossil cost against changes in the relative value of stored energy. The derivation and a verification argument, which shows that ρ^q does not depend on the initial state, are given in Appendix B. The HJB has no available closed form and is solved numerically as described in

Appendix F.[8]

The derivatives of W^q with respect to the storage levels give a direct economic characterization. When W^q is differentiable and neither storage nor power constraints bind, charging or discharging storage technology i equates marginal consumption value to the efficiency-adjusted shadow value of stored energy:

$$U'_\vartheta(c) = \begin{cases} \omega_{i+}\partial_{s_i} W^q, \alpha_i > 0, \\ \omega_{i-}\partial_{s_i} W^q, \alpha_i < 0. \end{cases} \tag{11}$$

If fossil production is interior, $U'_\vartheta(c) = p_f$. The shadow value $\partial_{s_i} W^q$ is thus the internal price of stored energy. Storage technology i is charged when the marginal value of consumption without its action would lie below $\omega_{i+}\partial_{s_i} W^q$, is discharged when it would lie above $\omega_{i-}\partial_{s_i} W^q$, and is idle in between. Thus, a technology with large conversion losses is idle for a larger shadow-price range. The shadow value is higher when future scarcity is more likely, for example in the scarcity regime or when storage levels are low, which gives storage the buffer-stock logic of precautionary saving (Deaton, 1991) and of competitive storage under uncertain production (Scheinkman and Schechtman, 1983).

2.6 Average Value, Capacity Costs, and the Optimal Mix

Capacity investment is evaluated using the optimal long-run hourly net operating benefit ρ^q from (7). The discounted present value is

$$\bar{V}^q = \int_0^\infty e^{-\delta t/T_Y} \rho^q dt = \frac{T_Y \rho^q}{\delta}. \tag{12}$$

Thus $T_Y \rho^q$ is the annual net operating benefit, or annual operating value, and $\bar{V}^q$ is its present value. The welfare-maximizing capacity mix therefore solves

$$q^* \in \arg\max_q \mathrm{NPV}^q, \qquad \mathrm{NPV}^q = \bar{V}^q - C(q), \tag{13}$$

where NPV^q is the net present value of the capacity mix and the cost term $C(q)$ in (13) contains fixed lifecycle capacity costs: upfront investment (CAPEX), refurbishment or replacement expenditure, and annual fixed operation and maintenance (O&M) costs.[9] Marginal costs of producing electricity with fossil fuel are part of $\bar{V}^q$. With unit lifecycle costs C_j, so that $C(q) = \sum_j C_j q_j$, an independently chosen positive capacity q_j at a differentiable interior optimum satisfies

[8] The method is a semi-Lagrangian scheme with an hourly time step, in which backward sweeps through the year are repeated until ρ^q is stable (relative value iteration).

[9] Output-dependent degradation is not modeled separately. The component-level lifecycle accounting, including the discounted refurbishment payments, is defined in Appendix D.

$$T_Y \partial_{q_j} \rho^{q^*} = \delta C_j. \tag{14}$$

The annual marginal operating benefit of each capacity thus equals its annualized lifecycle cost.[10] For a renewable capacity, the marginal operating benefit can be decomposed into the average output per unit of capacity times the average marginal value of electricity, plus the (endogenous) covariance between the two, so it falls when output is concentrated in states of surplus (Hirth, 2013). For storage energy capacity, it is the long-run average of the shadow value of stored energy times the net outflow from storage, per unit of capacity; this is positive because storage is filled when stored energy has a low shadow value and emptied when that value is high. Storage and fossil power capacities earn scarcity rents in the states in which their limits bind.

2.7 Consumption and Prices

At state $\xi = (x, r, s, \vartheta)$, let $c^q(\xi)$, $f^q(\xi)$, and $a_i^q(\xi)$ denote consumption, fossil dispatch, and the storage actions under the optimal stationary policy. The shadow price of electricity is the marginal consumption value

$$p^q(\xi) = U'_\vartheta(c^q(\xi)). \tag{15}$$

Below the curtailment threshold it exceeds p_{cap}, reflecting the marginal cost of unserved demand. A market with price cap p_{cap} would therefore have to ration demand in these states. Write $\bar{p}^q$ for the long-run average of the shadow price.

For a flow g_t under the selected optimal stationary policy for capacities q, write its long-run expected hourly average as

$$\mathcal{M}^q[g] = \lim_{T\to\infty} \frac{1}{T} \mathbb{E}_\xi \left[\int_0^T g_t dt \right]. \tag{16}$$

These averages are estimated from a large number of simulated operating paths. For a function h of the state, $\mathcal{M}^q[h]$ denotes the long-run average of the flow $h(\Xi_t)$; thus $\bar{p}^q = \mathcal{M}^q[p^q]$ and $\bar{c}^q = \mathcal{M}^q[c^q]$. The long-run standard deviation $\text{SD}(c^q)$ of consumption and the revenues and average sale prices of storage and renewable operation are defined in Appendix D.

Consumption can have very low marginal value while renewable production is abundant. Certainty-equivalent consumption $\bar{c}^q_{\text{CE}}$ is the annual average of the deterministic consumption profile, proportional to reference consumption, whose average benefit equals the long-run expected benefit of consumption.[11]

[10] Appendix C characterizes these marginal operating benefits, and the gradient-based optimization of capacities in Appendix F uses their discrete counterparts as the gradient of the NPV.

[11] The formal definition is given in Appendix D.

2.8 Levelized Cost of Electricity

A standard levelized cost of electricity (LCOE) divides discounted system cost by discounted expected consumption. Here system cost means fixed lifecycle capacity cost plus variable fossil cost; in a fully renewable system it is $C(q)$. Because large renewable capacities can produce high consumption when storage is full and marginal value is close to zero, equal standard LCOEs need not represent equal economic value when output differs in timing and value (Joskow, 2011; Ueckerdt et al., 2013).

Operating costs and consumption are evaluated as long-run hourly averages and discounted at the annual rate δ, as in the capacity criterion. Discounted system cost and discounted long-run consumption are

$$\mathrm{Cost}^q_{\mathrm{sys}} = C(q) + \frac{T_Y p_f}{\delta}\mathcal{M}^q[f^q], \quad \mathrm{Cons}^q = \frac{T_Y \bar{c}^q}{\delta}, \tag{17}$$

and discounted certainty-equivalent consumption $\mathrm{Cons}^q_{\mathrm{CE}}$ and discounted reference consumption $\mathrm{Cons}_{\mathrm{ref}}$ are defined in the same way from $\bar{c}^q_{\mathrm{CE}}$ and from the annual average $\bar{c}_{\mathrm{ref}}$ of reference consumption. The standard LCOE counts every MWh consumed equally, whereas the LCOE on certainty-equivalent consumption gives less weight to consumption of low marginal value:

$$\mathrm{LCOE}^q = \frac{\mathrm{Cost}^q_{\mathrm{sys}}}{\mathrm{Cons}^q}, \quad \mathrm{LCOE}^q_{\mathrm{CE}} = \frac{\mathrm{Cost}^q_{\mathrm{sys}}}{\mathrm{Cons}^q_{\mathrm{CE}}}. \tag{18}$$

I further introduce a welfare-adjusted LCOE that adds to system cost the monetary welfare loss from consumption below the phase-specific reference level and subtracts the welfare gain above it,

$$\mathrm{Deficit}^q = \frac{T_Y}{\delta}\mathcal{M}^q[\max\{-U_\vartheta(c^q), 0\}], \quad \mathrm{Surplus}^q = \frac{T_Y}{\delta}\mathcal{M}^q[\max\{U_\vartheta(c^q), 0\}]. \tag{19}$$

With the normalization $U_\vartheta(c_{\mathrm{ref}}(\vartheta)) = 0$, the adjustment $\mathrm{Deficit}^q - \mathrm{Surplus}^q$ is minus the discounted long-run consumption benefit, so the welfare-adjusted total cost is

$$\mathrm{Cost}^q_{\mathrm{tot}} = \mathrm{Cost}^q_{\mathrm{sys}} + \mathrm{Deficit}^q - \mathrm{Surplus}^q = C(q) - \bar{V}^q. \tag{20}$$

As the model is calibrated around c_{ref} and deviations are accounted for in the numerator, the denominator is taken to be discounted reference consumption. The welfare-adjusted reference-consumption LCOE is

$$\mathrm{LCOE}^q_{\mathrm{adj}} = \frac{\mathrm{Cost}^q_{\mathrm{tot}}}{\mathrm{Cons}_{\mathrm{ref}}} = -\frac{\mathrm{NPV}^q}{\mathrm{Cons}_{\mathrm{ref}}}. \tag{21}$$

While the fixed denominator is unusual, it ensures that minimizing the welfare-adjusted LCOE is equivalent to maximizing welfare.

This welfare-adjusted LCOE evaluates the mix as a whole and is unsuitable for evaluating one technology within the mix. The value-adjusted LCOE (VALCOE) (International Energy Agency, 2025) serves this purpose. It takes a marginalist approach: it evaluates an additional unit of a technology within a given mix, adjusting its LCOE for the value of its output relative to the system average, whereas the welfare-adjusted LCOE compares welfare levels across mixes.

3 Calibration

The dataset consists of German ENTSO-E production, load (consumption), and installed capacity time series from 2015 up to 2025 (European Network of Transmission System Operators for Electricity (ENTSO-E), 2026).[12] Except for installed capacities (yearly data), the data are quarter-hourly. Capacity factor time series are constructed by dividing renewable production by the interpolated installed capacity of the corresponding technology. The optimal mix is calibrated on onshore wind, in both capacity factors and costs (Table 2).[13] Because the capacity factors are computed from observed production, they include the historical effects of curtailment of wind and solar output, network constraints, and outages. The quarter-hourly capacity factors are averaged to hourly values, giving 96 432 hours. Observed day-ahead prices[14] serve to calibrate reference consumption and the reference price below and as a benchmark for the simulated prices in Section 5.2.

3.1 Wind and Solar Production

Both wind and solar production vary over the year and over the day. The daily cycle of solar production follows the position of the sun: production is zero at night and peaks around midday. The intraday solar profile is estimated directly from the hourly data in a separate step. The daily cycle of wind is weaker and mainly affects calm hours. The wind map is a seasonally shifted logistic function of a skewed transformation of the wind component.[15] The argument and the scale of the wind map also follow a daily cycle, so that calm hours cluster in the morning, as in the data. The solar map is a generalized Gamma quantile transformation of the scaled solar component, with seasonally varying shape and exponent.[16] All parameters other than the intraday solar profile and the bounds of the wind map, including those of the daily cycle of the wind map, the latent dynamics, and the drought regime, are estimated jointly by composite maximum likelihood (Varin, Reid and Firth, 2011) from pairs of hourly observations 24, 72, and 168 hours apart, with additional weight on pairs that end in hours of jointly low wind and solar

[12] The load and solar production series both cover only electricity exchanged with the public grid (Bundesnetzagentur, 2025). Electricity from rooftop photovoltaic installations that is consumed by their owners is included in neither series, so reference consumption is net of this self-consumption.

[13] The 2025 installed wind capacities shown for comparison cover onshore and offshore wind.

[14] These are the day-ahead prices of the German-Austrian-Luxembourg (DE-AT-LU) and German-Luxembourg (DE-LU) bidding zones over the same period, from the same source, averaged to 96 337 hourly values.

[15] The map is in the spirit of wind-speed models based on stochastic differential equations with non-Gaussian stationary distributions (Zarate-Minano, Mele and Milano, 2016). The logistic shape follows the power curve of a single turbine (its output as a function of wind speed), which Wang, Wang and Liu (2023) find to be well described by a logistic function, so the map can be read as the power curve of a representative turbine.

[16] This form is similar to positive stochastic irradiance models (Badosa et al., 2018).

production. The functional forms, the estimates, and the per-month model densities and data histograms are given in Appendix E.

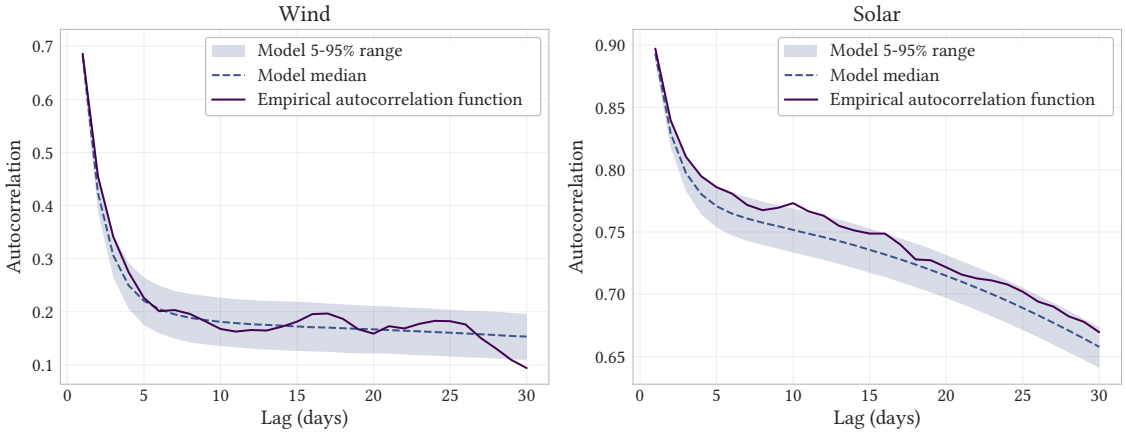


Figure 1: Autocorrelation functions of daily mean wind and solar capacity factors: data, and median and 5-95% range across simulated 11-year paths.

In Figure 1, the autocorrelation of daily mean capacity factors at a one-day lag is about 0.69 for wind and 0.90 for solar in the data, and the model gives 0.68 and 0.89. At a one-week lag, wind autocorrelation is 0.20 in the data and 0.19 in the model.[17] For solar, the model lies slightly below the data at lags beyond a few days. The model also matches the correlation between daily wind and solar output over the year, about −0.42. In winter, when diversification across the two technologies helps least, this correlation is only −0.10 in the data and −0.06 in the model.

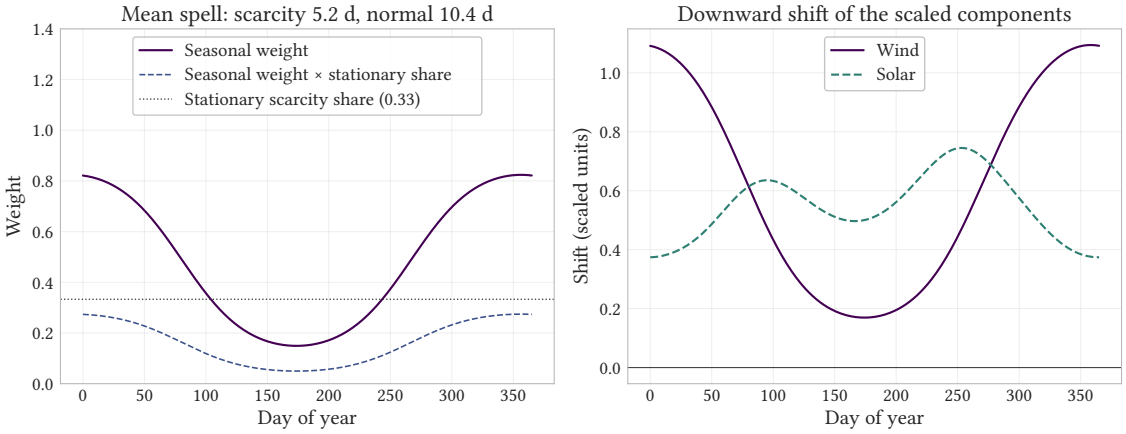


Figure 2: Seasonality of the drought regime: seasonal weight $D(\vartheta)/\bar{D}$ of the drought shift (see Appendix E), stationary scarcity share, and their product (left), and the downward shift of the scaled wind and solar components of the weather state at the drought mean of the latent state (right), with the wind scale evaluated without its daily cycle.

[17] These autocorrelations are computed from daily means that include the seasonal cycle, so they reflect both seasonal variation and the clustering of low- and high-production days.

Figure 2 describes the drought regime. Scarcity spells last 5.2 days and normal spells 10.4 days on average, so the scarcity regime is active about a third of the time. Its effect is strongly seasonal and concentrated in winter. The size of the shifts and the speed of adjustment after a switch are given in Appendix E. The regime thus generates periods of several days with jointly low wind and solar production in winter, which drive long-duration storage requirements. The model reproduces the share of days on which equal wind and solar capacities produce less than 30% of their mean output (3.1% simulated, 3.3% observed) and the share of winter days on which both daily capacity factors fall below their 10% quantiles (1.6% simulated, 1.5% observed). It somewhat understates the most extreme joint shortfalls and overstates the frequency of very calm wind days; see Appendix E. The model also reproduces the lower tail of 14- and 28-day mean capacity factors of wind-solar mixes in winter, the time scale of the shortfalls that hydrogen storage bridges (Figure 3, and Figure 23 in Appendix E).

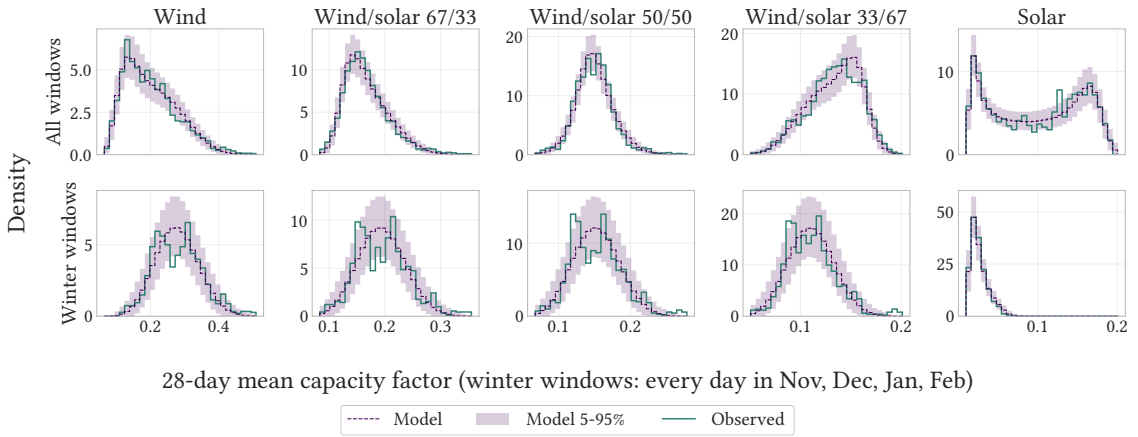


Figure 3: Densities of 28-day mean capacity factors for five wind-solar capacity mixes (columns), over all overlapping windows (top row) and over the windows that lie entirely in November-February (bottom row): observed (2015-2025), model, and the model's 5-95% range across 8 000 simulated paths of the same length.

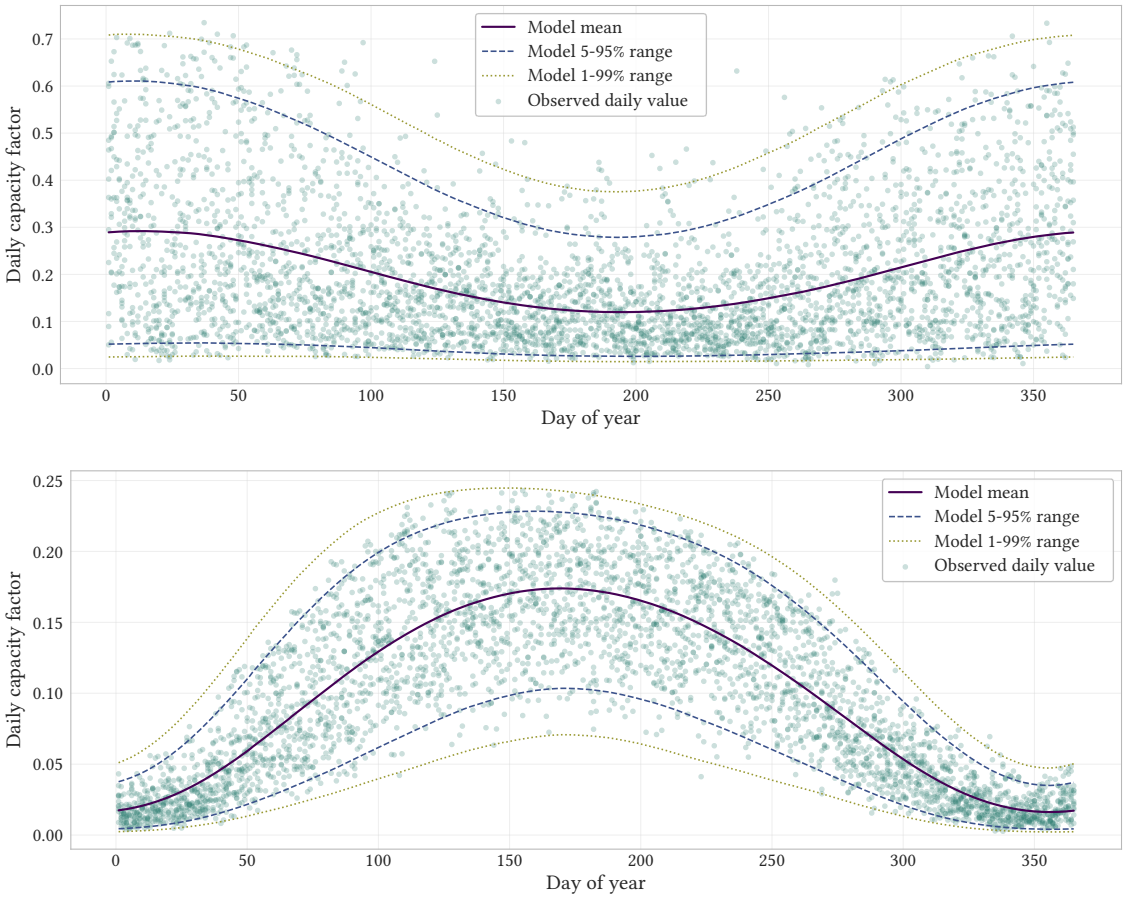


Figure 4: Daily mean wind (top) and solar (bottom) capacity factors over the year: model mean, 5-95% and 1-99% model quantiles, and observed daily values for 2015-2025.

In the top panel of Figure 4, wind output is higher and more dispersed in winter than in summer, and the observations share this broad seasonal pattern. The wide intervals within each season show why mean output alone is insufficient for sizing capacities: low-wind days remain possible even when expected generation is relatively high.

The corresponding solar curves in the bottom panel of Figure 4 reverse the seasonal pattern: output peaks in June and falls close to zero in winter. The observations support the timing and broad scale of this seasonal cycle. Together with the higher winter reference consumption shown below, the low winter solar mean raises the value of non-solar supply.

The estimated intraday solar profile varies seasonally and is normalized to a daily mean of one. Its construction is given in Appendix E. The model reproduces the timing of solar production over the day and its spread across days, and the weaker daily cycle of wind, with a dip in the morning and a rise in the afternoon. The hourly quantiles are compared with the data in Figure 24 and Figure 22 in Appendix E.

3.2 Reference Consumption

Reference consumption is derived from the observed German load in 2025, $c_{\text{obs}}(\vartheta)$, averaged to hourly values in UTC. The observed load of an hour is the quantity demanded at that hour's day-ahead price $p(\vartheta)$, whereas reference consumption is the quantity demanded at the reference price p_{ref} (Section 2). In hours priced above p_{ref}, the observed load is therefore moved along the calibrated inverse demand of Appendix A to the reference price,

$$c_{\text{ref}}(\vartheta) = c_{\text{obs}}(\vartheta)\left(\frac{\max\{p(\vartheta), p_{\text{ref}}\}}{p_{\text{ref}}}\right)^{|\varepsilon|}, \tag{22}$$

since all 2025 prices lie below the price cap p_{cap}. In hours priced at or below p_{ref}, reference consumption equals the observed load. The correction raises reference consumption in 4 374 hours, about half of the year, by at most 9.6%.[18] The same hourly profile, including the weekends and public holidays of 2025, is used in every simulated year and in every weather year of the perfect-foresight benchmark. Its mean is 53.5 GW, against 53.1 GW for the observed load, and hourly values range from 33.2 to 78.0 GW.

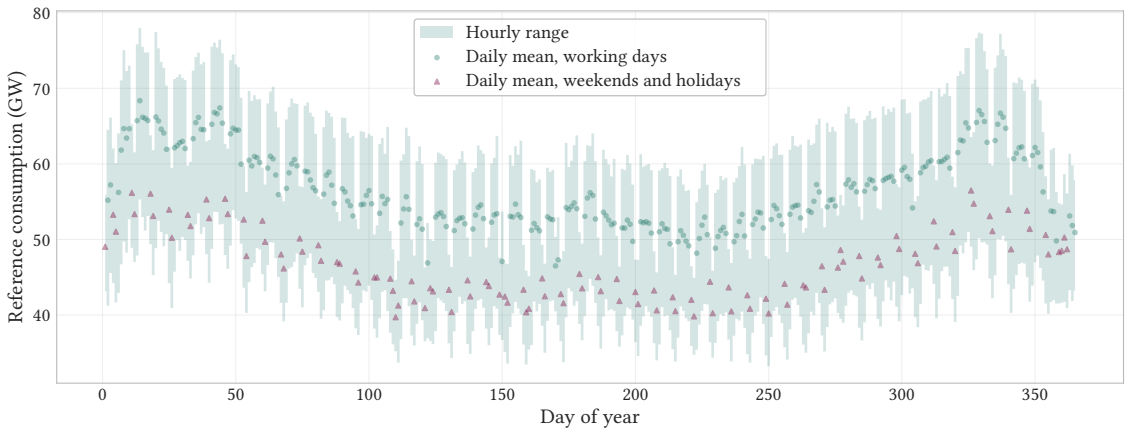


Figure 5: Daily mean reference consumption in 2025 over the year, with the range of hourly values.

In Figure 5, daily mean reference consumption is highest in January, reaching 68.4 GW, and lowest around Easter, falling to 39.7 GW, against 67.3 and 39.5 GW for the observed load. Weekends and holidays form a lower band throughout the year. Within the day, consumption is lowest overnight and rises steeply in the morning on working days (Figure 25 in Appendix E). These daily ramps create short-duration balancing needs, distinct from the slower seasonal movement in the daily mean.

[18] Appendix L gives the optimal mix without this correction and with a two-sided correction, which also lowers reference consumption in hours priced below p_{ref}.

3.3 Demand and Curtailment Costs

The monetary electricity consumption-benefit function specified in Appendix A is calibrated with the four parameters in Table 1.

Parameter	Symbol	Value	Basis
Blackout value	v_{blackout}	50 000 €/MWh	Assumption
Price cap	p_{cap}	5 000 €/MWh	Assumption
Reference price	p_{ref}	92.49 €/MWh	ENTSO-E, 2025
Elasticity	ε	-0.05	Short-run wholesale demand

Table 1: Parameters of the consumption-benefit function.

The maximum marginal value of electricity, interpreted as the marginal cost of unserved demand (at zero consumption), is assumed to be $v_{\text{blackout}} = 50\,000$ €/MWh[19]. The marginal-value threshold for demand curtailment (the price cap above which demand does not react to prices and demand curtailment takes place) is assumed to be $p_{\text{cap}} = 5\,000$ €/MWh[20]. The marginal value of electricity at reference consumption is set to the average of the hourly German day-ahead prices in 2025 weighted by the observed load, $p_{\text{ref}} = 92.49$ €/MWh, the average price paid per MWh in the year of reference consumption (European Network of Transmission System Operators for Electricity (ENTSO-E), 2026). The baseline elasticity is $\varepsilon = -0.05$, consistent with the short-run aggregate German wholesale-demand estimate in Hirth, Khanna and Ruhnau (2024). Appendix L reports the optimal mix under alternative blackout values and price caps.

This calibration implies $\kappa = k(\vartheta)/c_{\text{ref}}(\vartheta) = 0.819$, and hence a model-implied share of price-sensitive demand of 18.1%[21].

Figure 6 shows the consumption-benefit and marginal-value curves for a phase at elasticity $\varepsilon = -0.05$.

[19] Cambridge Economic Policy Associates (2018) estimate the value of lost load of German households at 12.41 €/kWh, that is, 12 410 €/MWh, with estimates for households across EU member states between 1.50 and 22.94 €/kWh and lower values for most industries. These estimates are averages over the consumption lost. In the model, the marginal value of consumption rises linearly from p_{cap} to v_{blackout} as consumption falls below the curtailment threshold, so the consumption lost below this threshold in a full blackout has an average value of $(p_{\text{cap}} + v_{\text{blackout}})/2 = 27\,500$ €/MWh.

[20] The threshold p_{cap} does not need to reflect actual market rules. For comparison, the harmonised maximum clearing price of the European single day-ahead market coupling has a reference value of 4 000 €/MWh and rises when clearing prices approach it (European Union Agency for the Cooperation of Energy Regulators, 2026).

[21] Hirth, Khanna and Ruhnau (2024) report that less than 25% of consumption is exposed to real-time prices and attribute the observed aggregate response primarily to industrial consumers. While the model-implied share of price-sensitive demand is consistent with this figure, the model is not separately calibrated on that share: it follows from the elasticity.

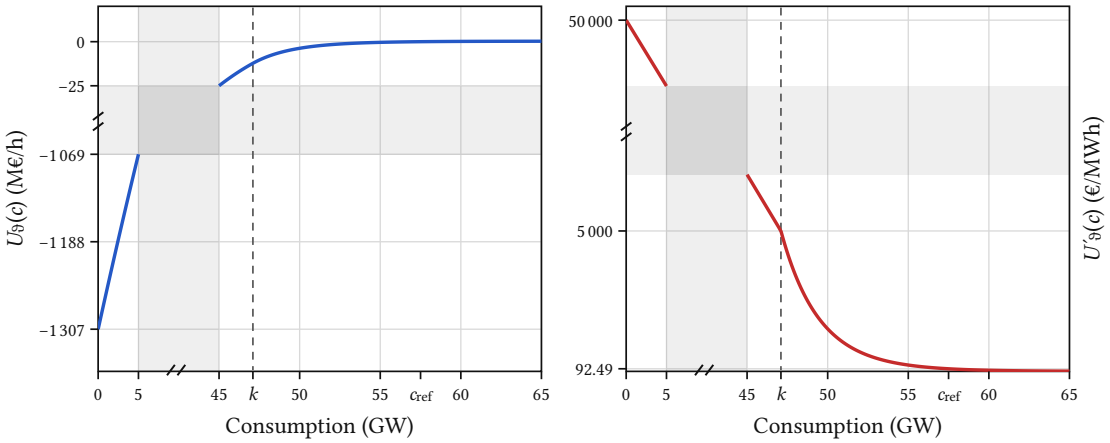


Figure 6: Monetary electricity consumption benefit and marginal value for a phase with $c_{\text{ref}} = 57.5$ GW.

3.4 Technology Costs and Efficiencies

The baseline cost assumptions for wind, solar, and battery storage are given in Table 2. For the baseline cases, these assumptions are from tables published by the Danish Energy Agency and use 2025 values except where only 2020 values were available (Danish Energy Agency, 2026a; 2026b). All costs are in 2025 euros.[22] The fixed O&M of wind includes its variable O&M.[23]

Variable	Technology	Data year	Unit	CAPEX [k€/unit]	Fixed O&M [k€/(unit · y)]	Lifetime [y]
K_W	Onshore wind	2025	MW	1 382	24.5	27
K_S	Utility-scale solar PV	2025	MW	595	10.0	40
$\bar{S}_B$	Battery storage	2025	GWh	350 140	10 154.1	15

Table 2: Technology-cost assumptions for wind, solar, and battery storage in the baseline cases.

The lifetimes enter the replacement-cost calculation, with a refurbishment fraction of 0.50 for every technology: at the end of each lifetime, half of the initial CAPEX is paid again to renew the installation (Appendix D).[24]

Batteries are assumed to be near-lossless, with charging efficiency $\omega_{B+} = \sqrt{0.91}$ and discharging efficiency parameter $\omega_{B-} = \frac{1}{\sqrt{0.91}}$, corresponding to a round-trip efficiency $\omega_{B+}/\omega_{B-} = 0.91$.

[22] Values that the Danish Energy Agency reports in 2020 euros are converted with the net price index it uses for its catalogues, a factor of 1.167 (Danish Energy Agency, 2026a). Battery costs are an exception: the Danish Energy Agency reports them in 2020 euros, deflated from 2024 prices with a factor of 0.84 based on the European Harmonised Index of Consumer Prices (HICP) (Danish Energy Agency, 2026b); I reverse this conversion and convert from 2024 to 2025 euros with the euro-area HICP, a factor of 1.021 (Eurostat, 2026a).

[23] The variable O&M is counted at 1 769 full-load hours per year, corresponding to the mean capacity factor of onshore wind in the observed 2015-2025 record.

[24] Appendix L reports the optimal mix with battery and hydrogen storage when the full initial CAPEX is paid at each renewal.

The hydrogen cost assumptions are listed in Table 3. They are taken from the Danish Energy Agency (Danish Energy Agency, 2026a; 2026c; 2026b). The hydrogen combined-cycle gas turbine (CCGT) has the fixed O&M and lifetime of a natural-gas CCGT and 1.15 times its CAPEX of 1 092 k€/MW.[25]

Variable	Technology	Data year	Unit	CAPEX [k€/unit]	Fixed O&M [k€/(unit · y)]	Lifetime [y]
$\bar{S}_H$	Hydrogen cavern storage	2020	GWh	3 723	79.2	100
K_{H+}	PEM electrolyzer	2025	MW	1 324	59.6	25
K_{H-}	Hydrogen CCGT	2020	MW	1 256	36.4	25

Table 3: Hydrogen technology-cost assumptions.

I assume that hydrogen storage has a charging efficiency (proton exchange membrane electrolyzers, PEM) of $\omega_{H+} = 0.60$ and a discharging efficiency parameter (hydrogen CCGTs) of $\omega_{H-} = 1/0.56 \approx 1.79$, corresponding to a CCGT efficiency of 0.56, the annual-average net efficiency of a natural-gas CCGT (Danish Energy Agency, 2026a). The round-trip efficiency is therefore 0.34. At a given shadow value of stored energy, the price above which hydrogen is discharged is thus $1/(0.60 \cdot 0.56) \approx 3.0$ times the price below which it is charged (Section 2), against $1/0.91 \approx 1.1$ for batteries.

[25] This is the CAPEX markup that Öberg, Odenberger and Johnsson (2022) assume for new gas turbines that burn pure hydrogen.

4 Results with Battery Storage Only

The battery-only case shows how a fully renewable mix compensates for the lack of cheap long-duration storage when the only storage technology has high round-trip efficiency but expensive energy capacity. This section restricts the model of Section 2 to batteries as the only storage technology, without fossil backup: $\bar{S}_H = K_{H-} = K_{H+} = F = 0$. Batteries are priced by energy capacity, and their power capacity allows a full charge or discharge in two hours, $K_{B+} = K_{B-} = \bar{S}_B/(2\,\mathrm{h})$, a constraint that is far from binding at the battery-only optimal energy capacities. The free capacities are therefore K_W, K_S, and $\bar{S}_B$.

The technology costs and efficiencies are given in Section 3 (Table 2).

The low-battery-cost case sets battery-storage CAPEX to 100 000 k€/GWh and fixed O&M to 2 900 k€/(GWh·y), 2.9% of CAPEX as in the baseline.[26]

The baseline combines the baseline technology costs with a 5% annual discount rate. The comparison cases use baseline technology costs at a 2% discount rate and low battery costs at a 5% discount rate. Wind, solar, and batteries have no fuel cost, and most of their annualized cost is the capital charge on investment. A lower discount rate therefore makes capacity cheaper relative to the annual welfare cost of shortages, and the optimum at 2% contains more of all three technologies (Table 4).

Case	Wind [GW]	Solar [GW]	Battery [GWh]
Baseline, 5%	547	601	728
Baseline, 2%	588	721	733
Low-battery-cost, 5%	435	341	2 628

Table 4: Optimized capacities with battery storage only.

Table 4 shows that the optimal wind-solar-battery mix has about seven and a half times the installed onshore and offshore wind capacity and seven times the installed solar capacity of 2025 at baseline battery costs, and its batteries alone can cover about 14 hours of average reference consumption. Low battery costs more than triple the optimal battery capacity, to about 49 hours of average consumption, and allow substantially lower wind and solar capacities. Batteries that hold 14 hours of average consumption cannot carry energy through scarcity spells, which last 5.2 days on average (Section 3). Supply during these spells can be raised only by installing enough wind and solar capacity that even low capacity factors deliver a large share of demand.

[26] If, for example, electric vehicles can be connected to the grid as storage or worn electric vehicle batteries can be reused as grid storage, battery storage may end up being substantially cheaper than cost estimates for new systems suggest.

An additional GW of wind or solar capacity is then valuable mainly in low-production hours and produces surplus of little value at other times. Cheaper batteries bridge longer shortfalls and thus replace part of this generation capacity. The cost elasticities at the battery-only optima show that wind, solar, and battery capacities substitute for one another (Table 24 in Appendix K).

Case	$\bar{c}_{\text{ref}}$ [GW]	$\bar{c}^q$ [GW]	SD (c^q) [GW]	$\bar{c}^q_{\text{CE}}$ [GW]	$\bar{p}^q$ [€/MWh]	$\delta\text{Cost}^q_{\text{sys}}$ [B€/y]	$\delta\text{Cost}^q_{\text{tot}}$ [B€/y]	LCOE [€/MWh]	LCOE_{CE} [€/MWh]	LCOE_{adj} [€/MWh]
Baseline, 5%	53.54	168.59	102.82	47.03	340.58	108.97	133.49	73.79	264.50	284.61
Baseline, 2%	53.54	188.69	117.47	47.89	249.66	81.29	98.04	49.18	193.78	209.04
Low-battery-cost, 5%	53.54	120.31	68.98	48.58	238.52	86.98	99.19	82.52	204.39	211.49

Table 5: Consumption, shadow price, system cost, welfare-adjusted cost, and LCOE summary with battery storage only. B€ denotes billions of euros.

Table 5 shows that arithmetic mean consumption is not a sufficient measure of the service delivered by these mixes. The gap between mean and certainty-equivalent consumption indicates that high-consumption outcomes raise the mean much more than they raise consumption valued through the concave consumption-benefit function.

Case	Curtailment $c \le k(\vartheta)$		Welfare loss $c \le k(\vartheta)$		Welfare loss $k(\vartheta) < c < c_{\text{ref}}(\vartheta)$		Total welfare loss		Welfare gain $c > c_{\text{ref}}(\vartheta)$	
	Hours [h/y]	Energy [TWh/y]	Annual [B€/y]	$/(T_Y\bar{c}_{\text{ref}})$ [€/MWh]	Annual [B€/y]	$/(T_Y\bar{c}_{\text{ref}})$ [€/MWh]	Annual [B€/y]	$/(T_Y\bar{c}_{\text{ref}})$ [€/MWh]	Annual [B€/y]	$/(T_Y\bar{c}_{\text{ref}})$ [€/MWh]
Baseline, 5%	169	1.623	25.02	53.34	1.56	3.33	26.58	56.67	2.07	4.41
Baseline, 2%	122	1.148	17.60	37.52	1.26	2.69	18.86	40.22	2.11	4.49
Low-battery-cost, 5%	98	0.791	12.11	25.82	2.06	4.39	14.17	30.21	1.95	4.16

Table 6: Curtailment and monetary welfare losses and gains relative to reference consumption with battery storage only. Curtailment: hours per year with $c \le k(\vartheta)$ and unserved energy below $k(\vartheta)$.

The decomposition in Table 6 locates most of the welfare loss on the curtailment branch $c \le k(\vartheta)$. Consumption lies on this branch in many hours each year, but the unserved energy, the shortfall of consumption below the curtailment threshold, remains a small fraction of reference consumption (469 TWh per year).

In Table 5, the standard LCOE spreads system cost over mean consumption, which is about three times reference consumption, although much of this consumption has little value. Both valuing consumption at its certainty equivalent (LCOE_{CE}, (18) in Section 2) and accounting for welfare losses and gains in the numerator with reference consumption as the denominator (LCOE_{adj}, (21) in Section 2) therefore raise the LCOE, from 74 to 265 and 285 €/MWh in the baseline case. The two adjusted measures differ by 3-8% across the three cases.

5 Results with Battery and Hydrogen Storage

The results extend the battery-only comparison by adding hydrogen storage and hydrogen conversion capacity. The cost and efficiency assumptions for hydrogen storage are given in Section 3 (Table 3).

Case	Wind [GW]	Solar [GW]	H_2 storage [GWh]	Battery [GWh]	Electrolyzer [GW]	H_2 CCGT [GW]
Baseline, 5%	312	216	35 010	44	47	45
Baseline, 2%	315	232	38 909	42	46	47
Low-battery-cost, 5%	271	249	40 047	314	30	41

Table 7: Optimized capacities with battery and hydrogen storage for the baseline and low-battery-cost cases.

The optimal mix in Table 7 features a large amount of hydrogen storage capacity and comparatively low battery storage capacity.[27] The optimal battery capacity is markedly lower than in the case where batteries are the only storage technology. This is due to the large difference in energy capacity costs: hydrogen storage provides a large stockpile of energy, able to insure against low-production events. Per GWh, cavern storage costs about 1% of battery energy capacity in CAPEX (3 723 against 350 140 k€).[28] The cost of energy capacity per MWh shifted falls with the number of storage cycles, and a stockpile turns over rarely: by the consumption shares in Figure 7, the batteries complete roughly 380 full cycles a year, whereas the hydrogen stock turns over only about three times. At the full 45 GW of hydrogen CCGT capacity, the hydrogen stock lasts about 430 hours, or 18 days, whereas batteries are emptied in about two hours. Wind and solar capacities produce much more electricity on average than reference consumption. This offsets low hydrogen round-trip efficiency and increases production in periods with low renewable capacity factors. Electrolyzers operate in periods of abundant renewable production, when electricity has a low shadow price, so the conversion losses fall on electricity of little value. Optimal renewable capacities are lower than with battery storage only (see Section 4), suggesting that raising production in periods with low capacity factors is a main reason for overbuilding renewables when batteries are the only storage, and that hydrogen storage takes over much of this role.

[27] These capacities lie far within the German resource base: the optimal 312 GW of wind capacity, and the 547 GW of the battery-only mix, compare with an onshore wind potential of about 1 190 GW under technical and ecological restrictions (Lütkehus, Salecker and Adlunger, 2013), and the 35 TWh of hydrogen storage with a technical potential of salt caverns in Germany of 9.45 PWh (Caglayan et al., 2020).

[28] Ziegler et al. (2019) estimate that meeting constant demand in every hour with wind, solar, and storage at a cost competitive with conventional generation requires storage energy capacity costs below about 20 USD/kWh. The cavern storage cost assumed here, about 3.7 €/kWh, lies well below this threshold, and the battery energy capacity cost, about 350 €/kWh, far above it.

Lower battery costs increase the optimal battery energy capacity substantially. The system also uses less wind and electrolyzer capacity and more hydrogen storage capacity, and solar capacity increases. More battery storage means that less energy has to make a round trip through inefficient hydrogen storage, and that, during the night, electrolyzers take up electricity stored in batteries, allowing the electrolyzers to operate more continuously in periods of abundant solar energy. Batteries also shift midday solar output into the evening and night, which raises the value of solar output and hence its optimal capacity. With battery storage only, low battery costs reduce solar capacity from 601 to 341 GW (Section 4), because batteries there replace the renewable overbuild that sustains production in low-production periods, a role that hydrogen storage takes over when both technologies are available. The increased hydrogen storage capacity reflects more interseasonal storage: while the storage remains too small for full interseasonal storage, more energy can be stored in summer and released in winter. Averaged over the simulated paths, the hydrogen storage level varies by about 20 TWh over the year with low battery costs, against about 10 TWh at baseline battery costs.

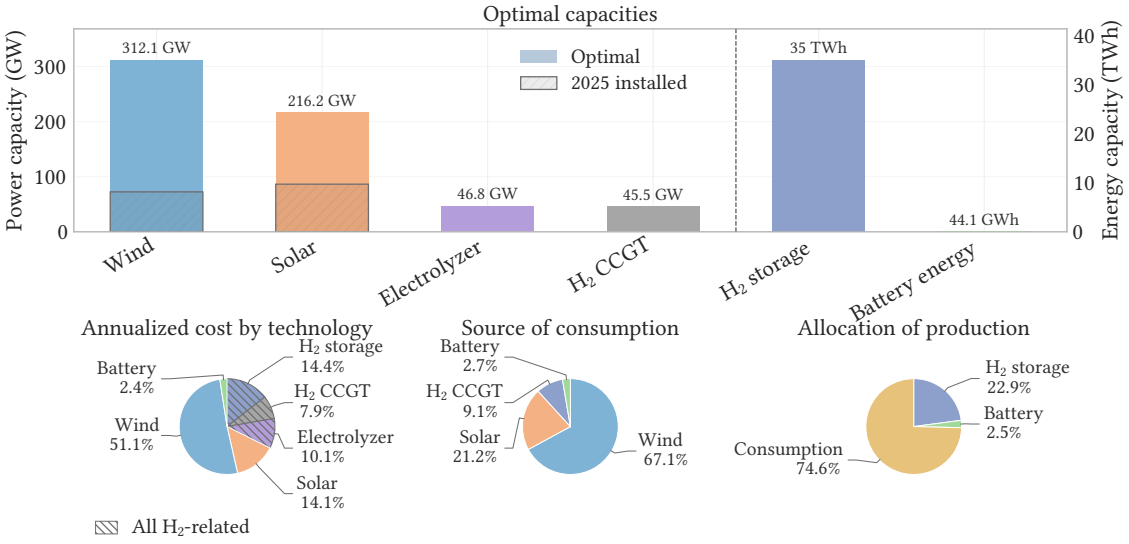


Figure 7: Optimal and 2025 installed capacities, annualized costs by technology, sources of consumed electricity, and allocation of produced electricity for the baseline case at a 5% discount rate. The 2025 installed wind capacity comprises onshore and offshore wind.

The consumption shares in Figure 7 and its low-battery-cost counterpart Figure 26 in Appendix I show that with low battery costs, batteries contribute more to consumption than hydrogen (10.1% against 6.2%), despite their much smaller energy capacity; at baseline battery costs, hydrogen contributes 9.1% and batteries 2.7%. Their repeated use for intraday balancing contrasts with hydrogen's role as a reserve for low-production periods. Consumption mostly comes directly from renewables because large capacities allow direct consumption much of the time.

The share is further inflated because excess production consumed with very low marginal value (and likely curtailed under real-world conditions) is counted as consumption.

Case	$\bar{c}_{\text{ref}}$ [GW]	$\bar{c}^q$ [GW]	SD (c^q) [GW]	$\bar{c}^q_{\text{CE}}$ [GW]	$\bar{p}^q$ [€/MWh]	$\delta\text{Cost}^q_{\text{sys}}$ [B€/y]	$\delta\text{Cost}^q_{\text{tot}}$ [B€/y]	LCOE [€/MWh]	LCOE_{CE} [€/MWh]	LCOE_{adj} [€/MWh]
Baseline, 5%	53.54	71.23	34.71	51.74	144.81	64.59	66.69	103.52	142.52	142.19
Baseline, 2%	53.54	73.23	35.81	52.89	98.06	44.75	45.35	69.75	96.58	96.69
Low-battery-cost, 5%	53.54	69.86	31.10	51.86	137.16	61.67	63.57	100.76	135.74	135.54

Table 8: Consumption, shadow price, system cost, welfare-adjusted cost, and LCOE summary with battery and hydrogen storage.

Case	Curtailment $c \le k(\vartheta)$		Welfare loss $c \le k(\vartheta)$		Welfare loss $k(\vartheta) < c < c_{\text{ref}}(\vartheta)$		Total welfare loss		Welfare gain $c > c_{\text{ref}}(\vartheta)$	
	Hours [h/y]	Energy [TWh/y]	Annual [B€/y]	$/(T_Y\bar{c}_{\text{ref}})$ [€/MWh]	Annual [B€/y]	$/(T_Y\bar{c}_{\text{ref}})$ [€/MWh]	Annual [B€/y]	$/(T_Y\bar{c}_{\text{ref}})$ [€/MWh]	Annual [B€/y]	$/(T_Y\bar{c}_{\text{ref}})$ [€/MWh]
Baseline, 5%	23	0.109	1.50	3.20	1.90	4.05	3.40	7.25	1.30	2.78
Baseline, 2%	13	0.060	0.83	1.77	1.19	2.54	2.02	4.31	1.42	3.03
Low-battery-cost, 5%	22	0.106	1.43	3.04	1.92	4.09	3.34	7.13	1.44	3.07

Table 9: Curtailment and monetary welfare losses and gains relative to reference consumption with battery and hydrogen storage. Curtailment: hours per year with $c \le k(\vartheta)$ and unserved energy below $k(\vartheta)$.

Monetary welfare losses relative to reference consumption on the curtailment branch, $c \le k(\vartheta)$, reflecting demand curtailment, account for somewhat less than half of the total welfare loss: such events are rare in the optimal mix, but each is costly. Consumption reaches the curtailment branch far less often than with batteries as the only storage technology, and unserved energy is more than an order of magnitude smaller (Section 4). In Table 9, most welfare losses arise from consumption between the curtailment threshold and the reference level. At the optimum, price-sensitive demand absorbs moderate shortfalls, which costs little compared to the conversion and storage capacities required to avoid them. Insurance against weather risk is therefore deliberately incomplete. Gains above reference consumption fall short of total losses in all three cases.

In Table 8, lowering the discount rate reduces all three LCOE measures much more than cheaper batteries do, highlighting the importance of the cost of capital. Relative to the battery-only mix (Table 5), hydrogen storage lowers the baseline welfare-adjusted LCOE from 285 to 142 €/MWh. The standard LCOE ranks the two mixes the other way, rising from 74 to 104 €/MWh, because the hydrogen mix produces far less low-value surplus consumption over which to spread its cost. The LCOE on certainty-equivalent consumption is close to the welfare-adjusted reference-consumption LCOE.

[29] Appendix I shows the same figures for the low-battery-cost case.

Figure 8, Figure 9, and Figure 10 illustrate the operation of the baseline optimal mix.[29] Along the simulated path, the hydrogen level falls from about 30 TWh in early January to about 6 TWh at the beginning of March and is rebuilt over the summer. This path reflects precautionary seasonal storage: the roughly 6 TWh left at the beginning of March is consistent with a buffer against a longer or deeper drought late in winter, since under perfect foresight the operator depletes the storage further (Section 6). In the December week, hydrogen discharge reaches the full 45 GW of CCGT capacity on most days, the battery runs empty on several days, and consumption stays just above the curtailment threshold. In the June week, which has low wind output, the battery is filled around midday and emptied in the evening, and on most days hydrogen is charged around midday and discharged during the night.

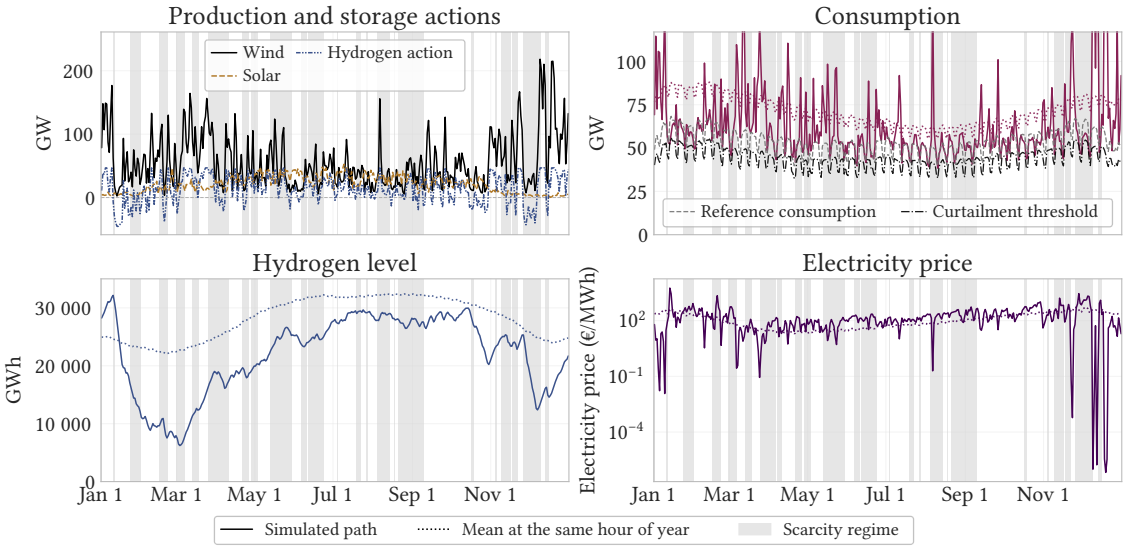


Figure 8: Simulated year beginning in January for the baseline case at a 5% discount rate: daily means along one simulated path, with the hydrogen storage level. The mean at the same hour of year is taken over all 4 096 simulated paths of 20 years each.

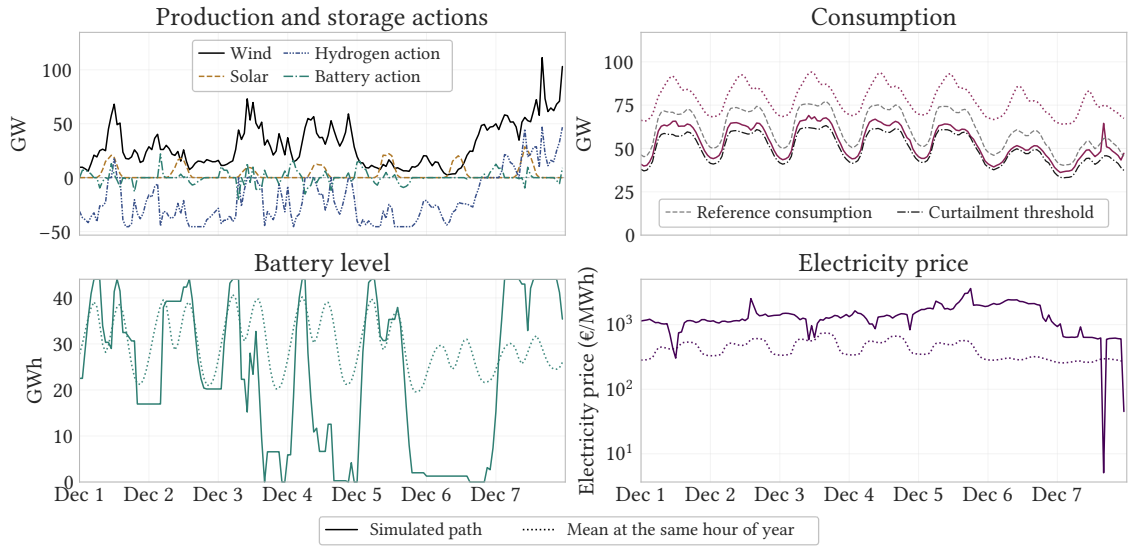


Figure 9: Simulated week in December for the baseline case at a 5% discount rate: hourly values along one simulated path, with the battery storage level. The mean at the same hour of year is taken over all 4 096 simulated paths of 20 years each.

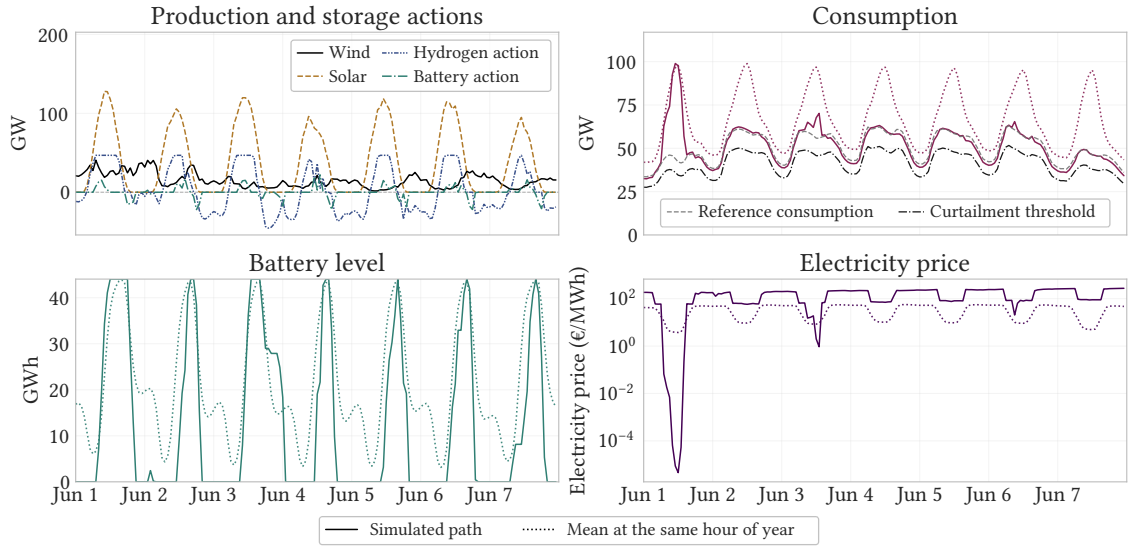


Figure 10: Simulated week in June for the baseline case at a 5% discount rate: hourly values along one simulated path, with the battery storage level. The mean at the same hour of year is taken over all 4 096 simulated paths of 20 years each.

5.1 Cost Elasticities

How the optimal mix responds to technology costs is summarized by the cost elasticities of the optimal capacities,

$$e_{ij} = \frac{\partial \ln q_i^*}{\partial \ln C_j}, \tag{23}$$

the percentage change in the optimal capacity q_i^* when the unit lifecycle cost C_j of capacity j rises by 1%, with all other unit costs fixed and the operation of the system adjusted to the new capacities. Appendix F derives them from the second derivatives of ρ^q at the optimum. Capacities i and j are substitutes if a higher cost of one raises the optimal amount of the other, and complements if it lowers it.

Table 10 reports the cost elasticities of the optimal mix with battery and hydrogen storage. Battery capacity is by far the most cost-elastic: its own-cost elasticity is −1.95 in the baseline case at a 5% discount rate, and it responds strongly to the costs of other technologies. A 1% higher unit cost of wind raises optimal battery capacity by 1.44%, and a 1% higher unit cost of electrolyzers by 0.68%, whereas higher costs of solar capacity and cavern storage lower it. Batteries thus substitute for wind and hydrogen conversion capacity, but complement hydrogen storage and solar capacity, whose midday output they shift into the evening and night. The own-cost elasticities of wind, solar, hydrogen storage, and electrolyzer capacity lie between −0.71 and −0.49. Wind substitutes for every other technology except the hydrogen CCGT: a 1% higher unit cost of wind raises optimal solar capacity by 0.71% and hydrogen storage capacity by 0.46%, so that a larger stockpile and more solar capacity make up for lower wind production in winter. Hydrogen CCGT capacity is almost insensitive to all unit costs, including its own (−0.10), consistent with its role of supplying the discharge power needed in the deepest shortfalls.

The elasticities at a 2% discount rate are close to those at 5%. With low battery costs, battery capacity is less cost-elastic (−1.10), whereas electrolyzer capacity responds more strongly to its own cost (−0.95) and becomes a complement to wind capacity and a substitute for solar capacity and batteries.

Case	Capacity	Unit cost of					
		Wind	Solar	H_2 storage	Battery	Electrolyzer	H_2 CCGT
Baseline, 5%	Wind	−0.494	0.196	0.129	0.068	0.056	−0.009
	Solar	0.709	−0.707	−0.066	−0.056	0.045	0.019
	H_2 storage	0.457	−0.065	−0.485	−0.044	−0.009	−0.006
	Battery	1.438	−0.326	−0.261	−1.947	0.678	0.338
	Electrolyzer	0.283	0.063	−0.013	0.163	−0.502	−0.041
	H_2 CCGT	−0.057	0.033	−0.011	0.104	−0.052	−0.098
Baseline, 2%	Wind	−0.494	0.193	0.115	0.075	0.070	−0.011
	Solar	0.699	−0.709	−0.045	−0.066	0.042	0.022
	H_2 storage	0.417	−0.045	−0.428	−0.042	−0.032	−0.004
	Battery	1.491	−0.365	−0.232	−2.094	0.734	0.388
	Electrolyzer	0.326	0.053	−0.042	0.171	−0.511	−0.045
	H_2 CCGT	−0.064	0.036	−0.007	0.116	−0.057	−0.096
Low-battery-cost, 5%	Wind	−0.714	0.341	0.265	0.116	−0.051	−0.007
	Solar	0.929	−0.895	−0.237	−0.125	0.240	0.024
	H_2 storage	0.711	−0.234	−0.596	−0.107	0.085	−0.002
	Battery	1.044	−0.411	−0.357	−1.101	0.654	0.105
	Electrolyzer	−0.355	0.608	0.219	0.503	−0.948	−0.061
	H_2 CCGT	−0.041	0.054	−0.004	0.073	−0.056	−0.107

Table 10: Cost elasticities e_{ij} of the optimal capacities with battery and hydrogen storage: the percentage change in the optimal capacity of the row when the unit lifecycle cost of the column rises by 1%.

5.2 Market Outcomes

With price-taking consumers and firms, the optimal allocation can be decentralized with the shadow prices as electricity prices (Appendix J), so the consumption, prices, and revenues below also describe market outcomes.[30]

Table 11 reports the quantity-weighted average sale prices (49) of Appendix D. In the baseline case, wind and solar sell at 59 and 48 €/MWh, well below the mean shadow price of 145 €/MWh, because their output is largest when electricity is abundant. Storage sells when electricity is scarce: net of charging expenditure, hydrogen storage earns 363 € and batteries 93 € per MWh discharged. With low battery costs, the larger battery capacity serves more shortfalls, and the net revenue of batteries falls to 50 € per MWh discharged. Hydrogen storage then discharges about a third less, mainly in deeper shortfalls, and earns 506 €/MWh.

[30]This requires uncapped marginal-value pricing, including on the curtailment branch. The effects of actual price caps are left for future research.

Case	Wind [€/MWh]	Solar [€/MWh]	H_2 storage [€/MWh]	Battery storage [€/MWh]
Baseline, 5%	59.01	47.62	362.85	92.55
Baseline, 2%	39.80	30.22	252.91	69.57
Low-battery-cost, 5%	59.39	47.65	505.96	49.62

Table 11: Quantity-weighted average sale prices with battery and hydrogen storage: wind and solar per MWh produced; storage: discharge revenue less charging expenditure, per MWh discharged. Hydrogen storage covers cavern storage, electrolyzers, and hydrogen CCGTs together.

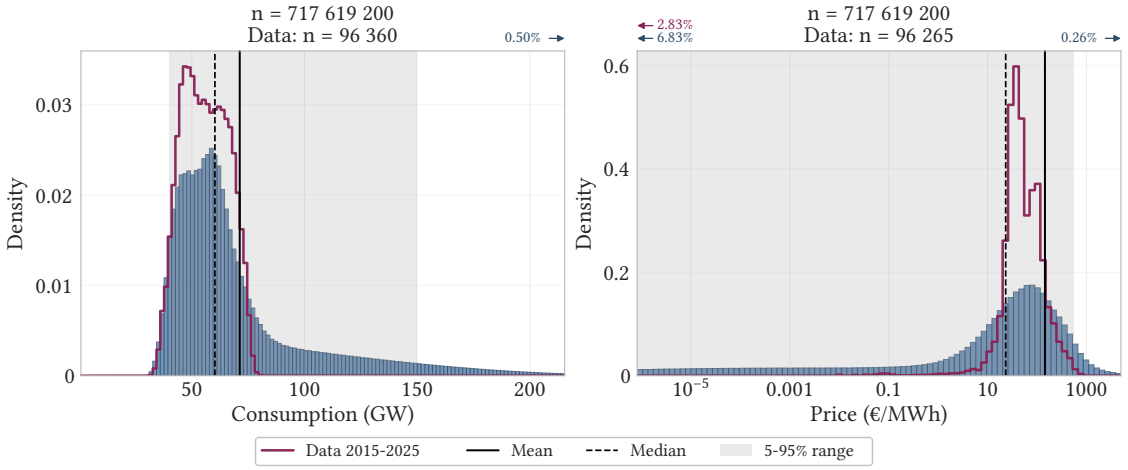


Figure 11: Simulated consumption and shadow-price histograms for the baseline case at a 5% discount rate, with the observed hourly load and day-ahead prices for 2015-2025 (Data). The mean, median, and 5-95% range refer to the simulated values; the arrows give the shares of simulated (lower arrows) and observed (upper arrows) values beyond the plotted range, and n counts hours.

Figure 11 and, for the low-battery-cost case, Figure 30 in Appendix I show right-skewed consumption distributions.[31] Relative to the observed hourly load in 2015-2025, much more probability is assigned to very high consumption, reflecting the treatment of surplus renewable production as consumption. Shadow prices span many orders of magnitude: electricity has almost no marginal value in surplus states, while a high-price tail remains for periods of low production and depleted storage. Lower battery costs reduce consumption dispersion and the mean shadow price, as quantified in Table 8, without removing this contrast.

[31] The simulated histograms pool the hours of 4 096 paths of 20 years each. The observed data exclude February 29, and the observed price series begins on 5 January 2015 and therefore has 95 fewer hours than the observed load. Nonpositive prices occur only in the observed data; they are omitted from the logarithmic price axis, and the upper arrow at the top left of the price panel reports their share among the observations.

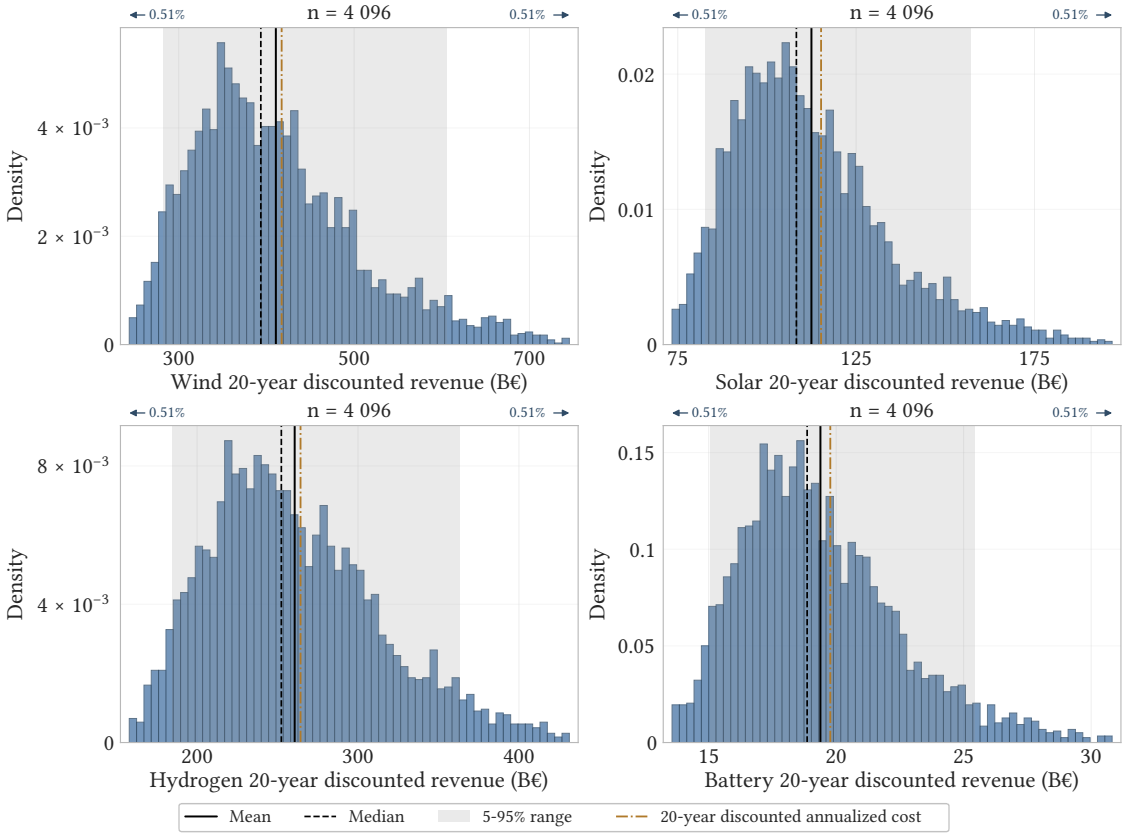


Figure 12: Simulated histograms of revenues over 20 years, discounted continuously at 5%, for the baseline case at a 5% discount rate. The dash-dotted line marks the annualized lifecycle capacity cost discounted over the same 20 years, the arrows give the shares of paths beyond the plotted range, and n counts the simulated paths. Hydrogen revenue and cost cover cavern storage, electrolyzers, and hydrogen CCGTs together.

The 20-year discounted revenue distributions in Figure 12 are right-skewed for all four technologies. Mean revenues are close to the discounted lifecycle costs, as the investment conditions (14) in Section 2 require at an interior optimum; Appendix J explains the small remaining gaps. Consumers and firms in the model are risk-neutral (Appendix J); if investors instead require compensation for this revenue risk, these findings suggest that a capacity mechanism may help avoid large risk premia.

5.3 Demand-Elasticity Sensitivity

More flexible demand may substitute for storage. This comparison holds the baseline technology costs and the 5% discount rate fixed, varying $|\varepsilon| \in \{0.03, 0.05, 0.10\}$; the case $|\varepsilon| = 0.05$ is the baseline optimum.[32]

[32]Increasing the elasticity magnitude $|\varepsilon|$ also increases the share of price-sensitive demand $1 - \kappa$ (Appendix A), which is 11.3%, 18.1%, and 32.9% for $|\varepsilon| = 0.03$, 0.05, and 0.10, respectively.

$\lvert\varepsilon\rvert$	Wind [GW]	Solar [GW]	H_2 storage [GWh]	Battery [GWh]	Electrolyzer [GW]	H_2 CCGT [GW]
0.03	319	216	39 383	44	48	49
0.05	312	216	35 010	44	47	45
0.10	300	217	25 875	43	42	38

Table 12: Optimized capacities for the demand-elasticity cases (baseline costs, 5% discount rate).

Across these cases (Table 12), increasing the elasticity magnitude reduces hydrogen storage from about 39 to 26 TWh, lowers electrolyzer and hydrogen CCGT capacities, and reduces wind capacity, while solar and battery capacities change little. These shifts are consistent with moderate shortfalls becoming cheaper to absorb, which lowers the value of winter supply from wind and hydrogen. Demand flexibility thus substitutes mainly for long-duration storage in the way it is modeled here.

$\lvert\varepsilon\rvert$	$\bar{c}_{\text{ref}}$ [GW]	$\bar{c}^q$ [GW]	SD (c^q) [GW]	$\bar{c}^q_{\text{CE}}$ [GW]	$\bar{p}^q$ [€/MWh]	$\delta\text{Cost}^q_{\text{sys}}$ [B€/y]	$\delta\text{Cost}^q_{\text{tot}}$ [B€/y]	LCOE [€/MWh]	LCOE_{CE} [€/MWh]	LCOE_{adj} [€/MWh]
0.03	53.54	72.41	36.24	52.00	145.05	67.10	69.21	105.77	147.30	147.56
0.05	53.54	71.23	34.71	51.74	144.81	64.59	66.69	103.52	142.52	142.19
0.10	53.54	69.93	33.43	51.35	143.05	59.45	61.65	97.04	132.17	131.45

Table 13: Consumption, shadow price, system cost, welfare-adjusted cost, and LCOE summary for the demand-elasticity cases (baseline costs, 5% discount rate).

$\lvert\varepsilon\rvert$	Welfare loss $c \le k(\vartheta)$		Welfare loss $k(\vartheta) < c < c_{\text{ref}}(\vartheta)$		Total welfare loss		Welfare gain $c > c_{\text{ref}}(\vartheta)$	
	Annual [B€/y]	$/(T_Y \bar{c}_{\text{ref}})$ [€/MWh]	Annual [B€/y]	$/(T_Y \bar{c}_{\text{ref}})$ [€/MWh]	Annual [B€/y]	$/(T_Y \bar{c}_{\text{ref}})$ [€/MWh]	Annual [B€/y]	$/(T_Y \bar{c}_{\text{ref}})$ [€/MWh]
0.03	1.79	3.81	1.16	2.48	2.95	6.29	0.84	1.78
0.05	1.50	3.20	1.90	4.05	3.40	7.25	1.30	2.78
0.10	0.96	2.04	3.58	7.63	4.54	9.67	2.33	4.97

Table 14: Monetary welfare losses and gains relative to reference consumption for the demand-elasticity cases (baseline costs, 5% discount rate).

Between $|\varepsilon| = 0.03$ and 0.10, the welfare-adjusted LCOE falls from 147.56 to 131.45 €/MWh and system cost by 7.7 B€/y (Table 13), while the net welfare loss rises slightly, from 2.1 to 2.2 B€/y (Table 14): curtailment losses fall and consumption above reference becomes more valuable, but losses between the curtailment threshold and reference consumption rise by more as a larger share of demand responds to prices. More elastic demand thus lowers the welfare-adjusted LCOE mainly through the lower capacity cost.

5.4 Robustness

Appendix L re-optimizes the baseline mix with a blackout value of 25 000 or 100 000 €/MWh, a price cap of 2 500 or 10 000 €/MWh, full payment of the initial CAPEX at each renewal, and

reference consumption without the price correction of Section 3 or with a two-sided correction. The welfare-adjusted LCOE lies between 141.12 and 143.40 €/MWh under the alternative blackout values and price caps and rises to 154.52 €/MWh under full renewal costs. Without the price correction, it is 141.79 €/MWh, and with the two-sided correction 144.59 €/MWh. In all seven cases, the optimal mix retains 305 to 313 GW of wind, 207 to 223 GW of solar, and between 33 and 37 TWh of hydrogen storage.

The results are insensitive to the size of the numerical grid, which supports the accuracy of the numerical method: doubling the number of nodes along every weather and storage dimension changes the NPV and the welfare-adjusted LCOE of the baseline optimum by about 0.15% (Appendix H). The computed solution also nearly satisfies the continuous-time HJB: along the stationary simulated paths, its residual averages 0.90 €/MWh of reference consumption, about 0.6% of the welfare-adjusted LCOE (Appendix F).

5.5 Hydrogen Storage Only

Removing batteries from the mix, $\bar{S}_B = K_{B+} = K_{B-} = 0$, shows how much they add once hydrogen storage is available. At baseline costs and a 5% discount rate, the optimal mix without batteries (Table 15) has more wind capacity than the mix with both storage technologies (326 rather than 312 GW), more electrolyzer and hydrogen CCGT capacity (51 GW each, rather than 47 and 45 GW), slightly less solar capacity, and about the same hydrogen storage capacity (34.4 rather than 35.0 TWh). Hydrogen conversion capacity thus takes over the intraday balancing that batteries provide in the mix with both technologies, and additional wind capacity makes up for the larger conversion losses. The welfare-adjusted LCOE rises only from 142.19 to 143.84 €/MWh at a 5% discount rate and from 96.69 to 97.78 €/MWh at 2% (Table 16), and curtailment and welfare losses change little (Table 17). Hydrogen storage then also discharges in hours that batteries would otherwise serve, and its net revenue per MWh discharged falls from 363 to 348 €/MWh at a 5% discount rate, while the average sale prices of wind and solar barely change (Table 23 in Appendix J).

Without batteries, the welfare-adjusted LCOE thus rises by about 1%, whereas without hydrogen storage it doubles (Section 4). This reinforces the importance of hydrogen storage relative to batteries: cheap long-duration storage determines the cost of a fully renewable mix, whereas batteries add little once it is available.

Case	Wind [GW]	Solar [GW]	H_2 storage [GWh]	Electrolyzer [GW]	H_2 CCGT [GW]
Baseline, 5%	326	210	34 427	51	51
Baseline, 2%	329	224	38 326	50	52

Table 15: Optimized capacities with hydrogen storage only.

Case	$\bar{c}_{\text{ref}}$ [GW]	$\bar{c}^q$ [GW]	SD (c^q) [GW]	$\bar{c}^q_{\text{CE}}$ [GW]	$\bar{p}^q$ [€/MWh]	$\delta\text{Cost}^q_{\text{sys}}$ [B€/y]	$\delta\text{Cost}^q_{\text{tot}}$ [B€/y]	LCOE [€/MWh]	LCOE_{CE} [€/MWh]	LCOE_{adj} [€/MWh]
Baseline, 5%	53.54	72.16	36.19	51.67	147.69	65.26	67.46	103.24	144.18	143.84
Baseline, 2%	53.54	74.12	37.16	52.82	99.96	45.19	45.86	69.60	97.66	97.78

Table 16: Consumption, shadow price, system cost, welfare-adjusted cost, and LCOE summary with hydrogen storage only.

Case	Curtailment $c \le k(\vartheta)$		Welfare loss $c \le k(\vartheta)$		Welfare loss $k(\vartheta) < c < c_{\text{ref}}(\vartheta)$		Total welfare loss		Welfare gain $c > c_{\text{ref}}(\vartheta)$	
	Hours [h/y]	Energy [TWh/y]	Annual [B€/y]	$/(T_Y \bar{c}_{\text{ref}})$ [€/MWh]	Annual [B€/y]	$/(T_Y \bar{c}_{\text{ref}})$ [€/MWh]	Annual [B€/y]	$/(T_Y \bar{c}_{\text{ref}})$ [€/MWh]	Annual [B€/y]	$/(T_Y \bar{c}_{\text{ref}})$ [€/MWh]
Baseline, 5%	26	0.115	1.60	3.42	1.88	4.01	3.48	7.43	1.28	2.74
Baseline, 2%	15	0.063	0.90	1.91	1.18	2.52	2.08	4.43	1.41	3.00

Table 17: Curtailment and monetary welfare losses and gains relative to reference consumption with hydrogen storage only. Curtailment: hours per year with $c \le k(\vartheta)$ and unserved energy below $k(\vartheta)$.

6 The Role of Uncertainty

Capacity models that optimize dispatch with perfect foresight on historical weather paths assume that the operator knows future renewable production. This section asks how much this assumption changes the operation of the optimal mix, its operating value, and the optimal capacities. I compare the expected operating value that the optimal mix (see Table 7) attains under the optimal policy of the stochastic problem with the value that the same mix attains when the weather path is known in advance (see Appendix M for the problem statement and further details).

Figure 13, Figure 14, and Figure 15 compare the operation of this mix under weather uncertainty and under perfect foresight.[33] Optimal operation under perfect foresight anticipates future production and depletes the storage further. Under uncertainty, the operator does not know when the next surplus will arrive and keeps a higher stock as a precaution. The storage is therefore more often full when a surplus arrives, and the surplus that cannot be stored is consumed at extremely low marginal value, which raises average consumption. The consumption benefit is lower, the shadow price is higher, and both vary more under uncertainty. In absolute terms, the relevance of perfect foresight for the operating benefit is particularly pronounced in winter, when production is dominated by persistent, uncertain wind energy.

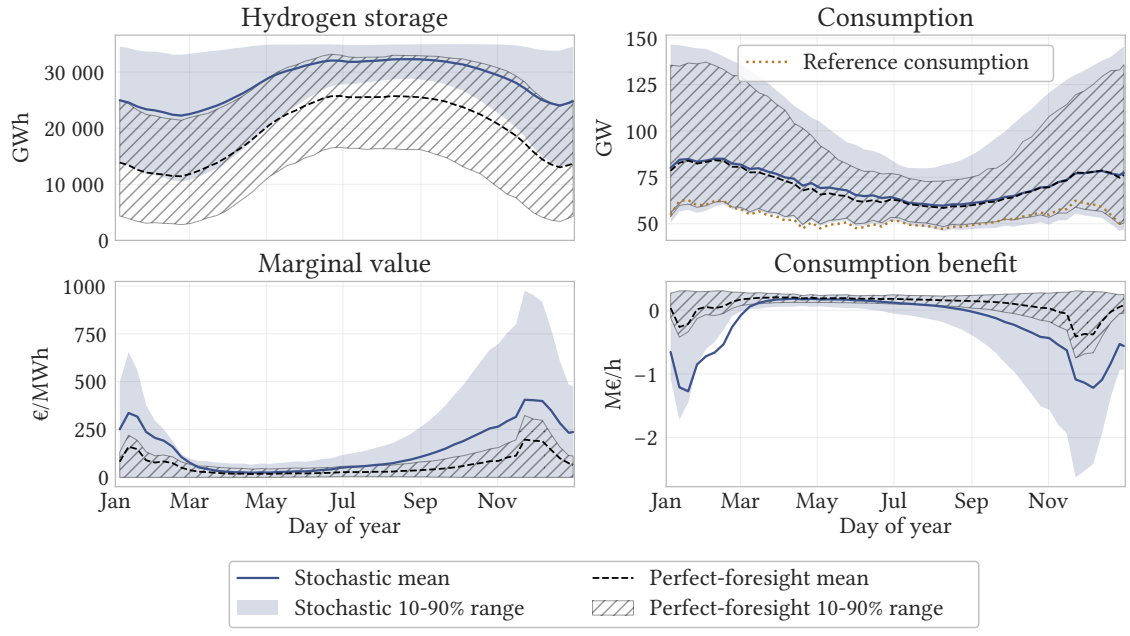


Figure 13: Stochastic and perfect-foresight operation over the year: weekly averages of the hourly simulation means and 10-90% ranges for the baseline mix at a 5% discount rate.

[33] The perfect-foresight problem is solved with an infinite horizon on 256 simulated 10-year paths, each repeated periodically, giving 2 560 years in the figures.

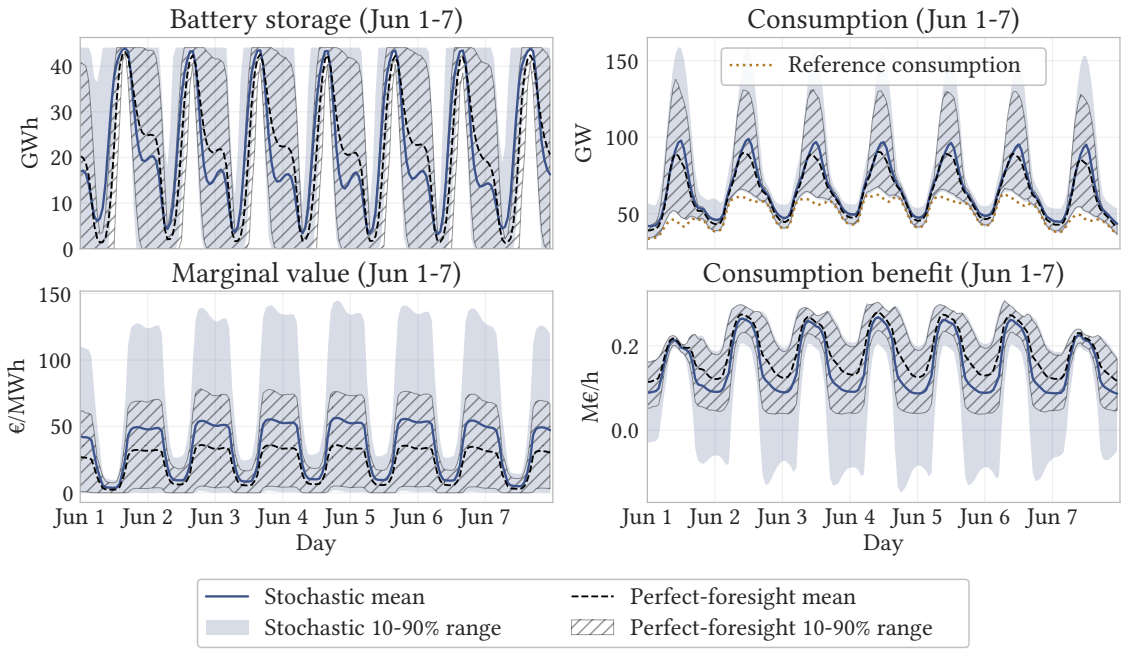


Figure 14: Stochastic and perfect-foresight operation, June 1-7: simulation means and 10-90% ranges for the baseline mix at a 5% discount rate.

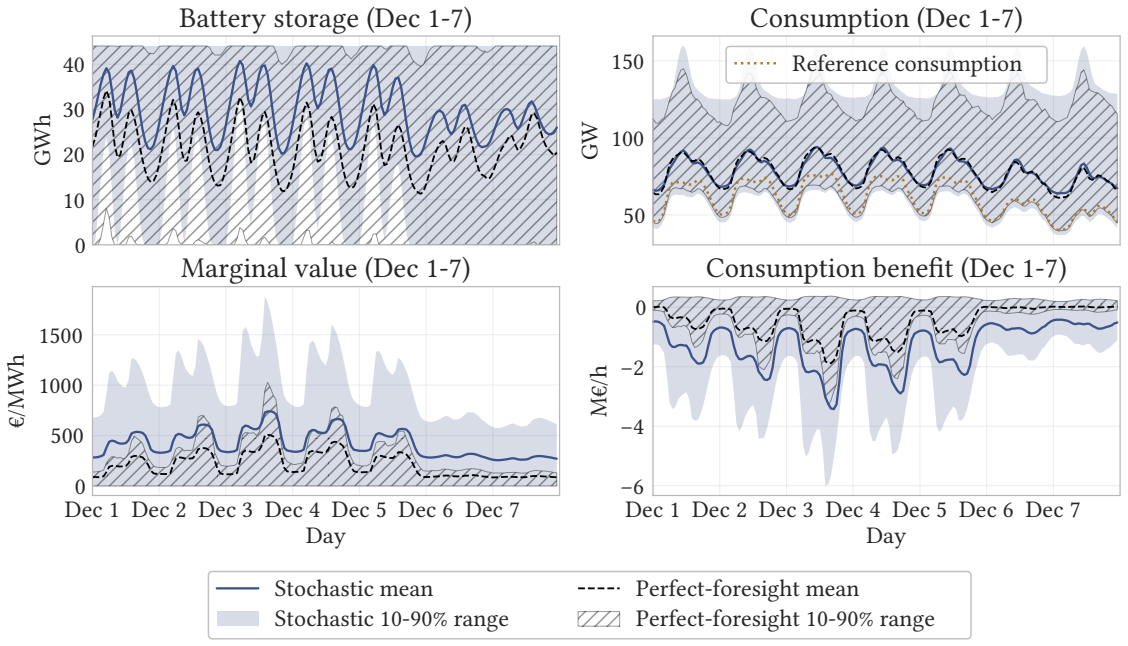


Figure 15: Stochastic and perfect-foresight operation, December 1-7: simulation means and 10-90% ranges for the baseline mix at a 5% discount rate.

Yet, differences in the optimal operation do not imply that the operating value needs to be vastly different when compared to system costs. Figure 16 shows this comparison for the optimal mix under uncertainty (left violin) and for the 2016-2025 perfect-foresight optimum, the mix that is optimal when the weather of the observed decade is known in advance (right violin). Recall

that the consumption-benefit function is normalized and assigns zero operating value when consumption equals reference consumption.

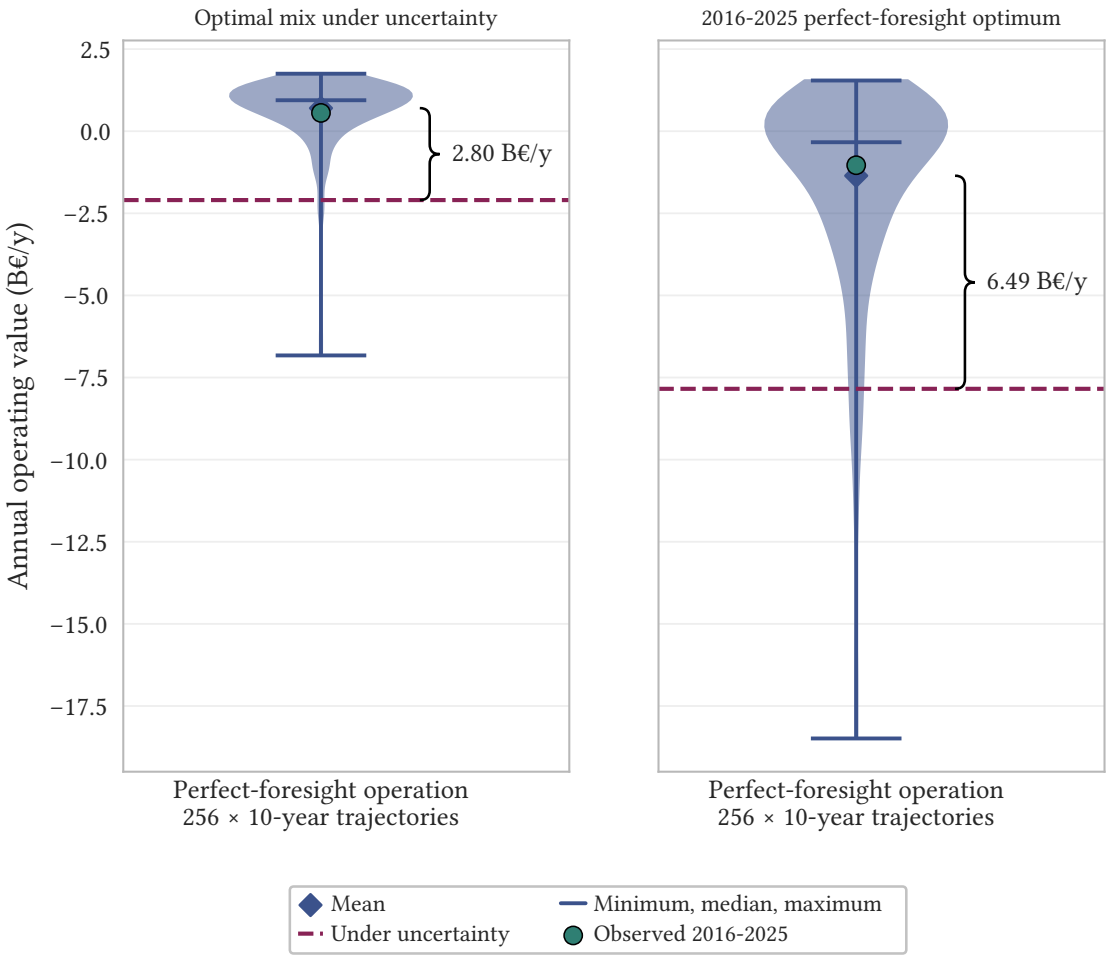


Figure 16: Annual operating values, under perfect foresight and under uncertainty (dashed line), of the optimal mix under uncertainty (left) and of the 2016-2025 perfect-foresight optimum (right) across the 256 simulated 10-year paths; the number between the mean and the dashed line gives the average value of perfect foresight of that mix. The circle marks the perfect-foresight operating value of each mix on the observed 2016-2025 data.

With $\rho^q_{\mathcal{P}_n}$ denoting the perfect-foresight average hourly operating value (72) in Appendix M on the n-th of the $N = 256$ simulated 10-year paths, the average value of perfect foresight of a mix q is $\frac{1}{N}\sum_{n=1}^{N} T_Y \rho^q_{\mathcal{P}_n} - T_Y \rho^q$. Both terms depend on the mix. In Figure 16, this average value is the distance between the mean of the perfect-foresight values and the dashed line marking $T_Y \rho^q$ for each violin, annotated with its value for that mix.

For the optimal mix under uncertainty (left violin), the average value of perfect foresight is around 2.8 B€/y, compared to annualized system costs of around 65 B€/y. Depending on the realization of the 10-year path, however, its perfect-foresight operating value can vary by more

than 8.5 B€/y, from about −7 to 2 B€/y. The 2016-2025 perfect-foresight optimum (right violin) holds less hydrogen storage (Table 18). Its perfect-foresight operating value varies much more across the simulated paths, from about −18 to 2 B€/y, and its average value of perfect foresight is around 6.5 B€/y, more than twice that of the optimal mix under uncertainty.

Under weather uncertainty, the expected operating value of the 2016-2025 perfect-foresight optimum is around −7.8 B€/y, about 5.8 B€/y below the value of around −2.1 B€/y for the optimal mix under uncertainty. This loss outweighs its capacity-cost saving of only about 3.1 B€/y, so the optimal mix under uncertainty performs better by about 2.7 B€/y. The comparison of the two mixes suggests that information and stored energy are substitutes. The optimal mix under uncertainty holds a large buffer, affordable because cavern storage capacity is cheap, so knowing the weather in advance adds little to its operating value. The 2016-2025 perfect-foresight optimum holds a smaller stock, which makes perfect foresight more valuable to it and weather uncertainty more costly. The low value of perfect foresight thus relies on cheap long-duration storage and would likely be higher if such storage were expensive.[34]

This leads to the question of whether the optimal mix differs materially when the weather is known in advance. Table 18 summarizes the perfect-foresight optima at a 5% discount rate for the observed decade, its two halves, and individual (periodically repeated) years.[35]

Period	Wind [GW]	Solar [GW]	H_2 storage [GWh]	Battery [GWh]	Electrolyzer [GW]	H_2 CCGT [GW]	Capacity cost [B€/y]	NPV [T€]	LCOE [€/MWh]	$LCOE_{CE}$ [€/MWh]	$LCOE_{adj}$ [€/MWh]
2017	304	196	21 313	37	50	48	59.69	−1.21	105.32	129.59	129.26
2019	270	199	15 387	58	39	46	53.73	−1.08	92.39	115.75	115.61
2021	336	180	17 348	38	48	40	60.35	−1.22	103.82	130.78	130.44
2023	257	196	17 942	37	47	46	53.21	−1.07	95.99	114.52	114.40
2025	299	225	22 258	48	45	46	60.16	−1.22	110.10	130.30	129.96
2016-20	326	213	23 422	32	49	45	62.66	−1.27	99.34	136.34	135.89
2021-25	301	201	22 437	34	49	46	59.41	−1.21	102.84	128.92	128.60
2016-25	317	208	23 263	31	49	46	61.52	−1.25	100.65	133.74	133.34
Uncertainty	*312*	*216*	*35 010*	*44*	*47*	*45*	*64.59*	*−1.33*	*103.52*	*142.52*	*142.19*

Table 18: Perfect-foresight optima at a 5% discount rate for the observed decade, its two halves, and five individual years, and the optimal mix under uncertainty (italics) for comparison. T€ denotes trillions of euros.

For the observed 2016-2025 decade, the perfect-foresight mix has slightly more wind and slightly less solar capacity than the optimal mix under uncertainty. It holds much less hydrogen storage (23 rather than 35 TWh) and less battery capacity, with similar electrolyzer and hydrogen CCGT capacities. The two optima differ mainly in the stockpile held as insurance against low-production periods. For a given storage capacity, Geske and Green (2020) find that the value of

[34] An operator who does not observe the drought regime chooses nearly the same mix, at a slightly lower NPV (Appendix N).

[35] Further results are provided in Appendix M.

storage under perfect foresight is 27% higher than under uncertainty, because perfect foresight allows arbitrage that precautionary storage would forgo. When storage capacity is chosen as well, perfect foresight nevertheless leads to a smaller stockpile. Strictly speaking, this comparison does not, however, isolate the value of knowing the weather in advance: it also reflects differences between the observed decade, repeated cyclically, and the weather distribution of the calibrated model. On the 2016-2025 data, the perfect-foresight operating value of each mix lies within the range of the simulated paths (with the model calibrated mostly on this decade), close to their mean (Figure 16): slightly below it for the optimal mix under uncertainty and slightly above it for the 2016-2025 perfect-foresight optimum.

The individual years instead imply optima that differ substantially from one another. Single-year paths are too short: the optima they imply reflect the weather of that year rather than the distribution of weather that a capacity decision has to cover.

7 Results with Fossil Backup

This section allows natural-gas backup at increasing carbon prices, interpreted as the social cost of carbon, to distinguish the cost of a zero-emission system from the marginal cost of eliminating the remaining fossil generation. It asks at which carbon price hydrogen storage enters the optimal mix and how the abatement cost rises as emissions approach zero.

7.1 Model Specifications and Assumptions

The hydrogen CCGTs used for discharging hydrogen storage are hydrogen-ready: they can also burn natural gas, whereas hydrogen can only be burned in these hydrogen-ready CCGTs. Thus,

$$0 \leq f_t, \qquad -a_{Ht-} \leq K_{H-}, \qquad f_t - a_{Ht-} \leq F + K_{H-}. \tag{24}$$

For fossil-only CCGT capacity F, I use a CAPEX of 1 092 k€/MW, fixed O&M of 36.4 k€/(MW·y), and a lifetime of 25 years (Danish Energy Agency, 2026a). The refurbishment fraction remains at 0.50. The discount rate is fixed at 5% in this section.

I assume that the emissions intensity of natural-gas CCGT γ_f is given by

$$\gamma_f = 0.36 \text{ tCO}_2/\text{MWh}, \tag{25}$$

the emission factor of natural gas in Germany, 55.83 tCO_2/TJ of fuel (Juhrich, 2022), at the CCGT efficiency of 0.56. Then, for a carbon price τ in €/tCO_2, interpreted as the social cost of carbon, and a marginal cost of producing electricity with natural gas (including fuel and marginal operating costs) of 108.10 €/MWh, the total marginal social cost of producing electricity with natural gas is given by

$$p_f = 108.10 + \gamma_f \tau. \tag{26}$$

This marginal cost corresponds to the EU natural-gas price of July 2026.[36] Appendix O reports the lifetime cost of supplying reference consumption with natural gas, hard coal, or lignite alone.

Annual fossil emissions in ktCO_2/y under the optimal stationary policy are

$$E^q = \gamma_f T_Y \mathcal{M}^q[f^q]. \tag{27}$$

7.2 Capacities

ENTSO-E's 2025 German generation capacity data and the optimum at zero carbon price are compared in Table 19 (Aichner et al., no date; European Network of Transmission System

[36] The price is converted to 2025 euros with the euro-area HICP of July 2026 (Eurostat, 2026b) and from the gross to the net calorific value of natural gas (by the factor 1.108 of Juhrich (2022)), divided by the CCGT efficiency, and raised by variable O&M of 5.46 €/MWh.

Operators for Electricity (ENTSO-E), 2026).[37]

Technology	Unit	Germany 2025	Optimum at zero carbon price
Wind	GW	72	141
Solar	GW	86	124
Gas-capable generation	GW	37	55
of which hydrogen CCGT	GW	—	—
of which fossil-only	GW	37	55
Electrolyzer	GW	—	—
Battery power	GW	11	6 (implied)
Battery energy	GWh	25	11
H_2 storage	GWh	—	—
Coal	GW	31	—
Pumped-storage power	GW	9	—
Pumped-storage energy	GWh	36	—

Table 19: Installed German capacities in 2025 and optimal capacities at zero carbon price. A dash in the installed column marks unavailable data and presumably negligible capacity; a dash in the optimum column marks a technology outside the optimal mix, which, except for hydrogen storage and conversion, reflects a model restriction.

The zero-carbon-price optimum has substantially more renewable capacity than the 2025 system. It also has more gas-fired capacity (due to the absence of coal in the model) and less than half the battery energy capacity. Renewable expansion is therefore worthwhile under the assumed costs even without pricing emissions, but does not eliminate dispatchable gas generation. Without a carbon price, natural gas sets the shadow price at $p_f = 108.10$ €/MWh whenever it is marginal, so renewable output is worth the fuel it saves in those hours and little in surplus hours; wind and solar capacities expand until this value just covers their cost. Because natural gas caps the shadow price at p_f whenever its capacity does not bind, the price spreads available to storage are small in most hours, which supports a small battery capacity and no hydrogen storage.

[37] Installed battery power is ENTSO-E's non-pumped "Energy storage" category, installed gas covers all gas generation, and installed wind comprises 63 GW onshore and 9 GW offshore; the optimum uses onshore wind and gas-capable CCGT only.

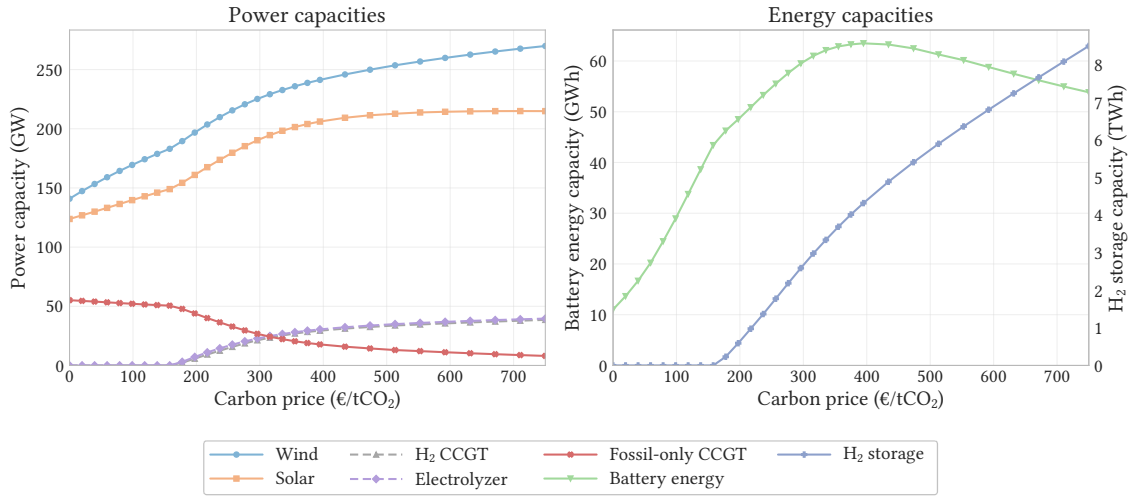


Figure 17: Optimized power and energy capacities against the carbon price.

In Figure 17, hydrogen storage first enters the optimal mix at a carbon price between 158 and 178 €/tCO₂. Below this carbon price, emissions fall through more wind, solar, and battery capacity: between a zero carbon price and 158 €/tCO₂, wind capacity rises from 141 to 183 GW, solar capacity from 124 to 149 GW, and battery capacity from 11 to 43 GWh, while fossil-only capacity falls only from 55 to 51 GW. Battery energy capacity keeps rising after hydrogen storage enters, to a peak of about 63 GWh near 395 €/tCO₂, and then falls to 54 GWh at 750 €/tCO₂. Batteries charge from midday solar output and discharge in the evening, so they displace natural gas on many days, whereas hydrogen storage covers rarer and longer shortfalls. At high carbon prices, electrolyzers absorb more of the midday surplus and hydrogen CCGTs cover more of the evening demand, as a fossil backup does at lower carbon prices. Solar capacity levels off at about 215 GW from about 550 €/tCO₂, while wind capacity keeps rising. At a carbon price of about 300 €/tCO₂, within the range of recent estimates of the social cost of carbon (Rennert et al., 2022), the optimal mix has 225 GW of wind, 190 GW of solar, 60 GWh of battery storage, 2.6 TWh of hydrogen storage with 23 GW of electrolyzer and 21 GW of hydrogen CCGT capacity, and 27 GW of fossil-only capacity. It emits 19 MtCO₂/y, about 95% below the 1990 emissions of the German electricity mix, at a welfare-adjusted levelized cost of 116 €/MWh.

At 750 €/tCO₂, hydrogen storage capacity remains well below that of the zero-emission optimum without fossil backup (Table 7). Dedicated fossil-only capacity falls to about 8 GW at the highest carbon price considered, 750 €/tCO₂, where hydrogen-ready CCGTs still burn natural gas. Since hydrogen-ready CCGTs burn both fuels, the choice is between natural gas, at the marginal cost p_f and without additional capacity, and stored hydrogen, which requires electrolyzer and cavern capacity but is produced from surplus electricity of low value. Hydrogen is therefore cheaper for energy needed often, over which its capacity cost is spread, and natural

gas for energy needed rarely. As the carbon price raises p_f, hydrogen becomes the cheaper option at lower frequencies of use, so its capacities expand. The last MWh of natural gas, however, is used only in the worst droughts. At a 5% discount rate, a GWh of cavern capacity costs about 266 k€ per year (Table 3), so the cavern alone costs about 475 €/MWh of electricity if it is filled and emptied once a year, more than the 378.1 €/MWh of natural gas at 750 €/tCO_2, and ten times as much if it is used once a decade. Natural gas therefore remains in the mix even at the highest carbon price considered.

7.3 Abatement Cost

For a carbon price τ, let q^*_τ denote the optimal capacity vector, with $\bar{V}^{q^*_\tau}$ and $E^{q^*_\tau}$ evaluated under the corresponding optimal stationary policy. Relative to a zero carbon price, the average abatement cost in Figure 18 is

$$\mathrm{AC}(\tau) = \frac{\delta(C(q^*_\tau) - C(q^*_0)) - \left(\delta \bar{V}^{q^*_\tau} + \tau E^{q^*_\tau} - \delta \bar{V}^{q^*_0}\right)}{E^{q^*_0} - E^{q^*_\tau}}, \tag{28}$$

where $\delta C(q)$ is annualized capacity cost and $\delta \bar{V}^{q^*_\tau} + \tau E^{q^*_\tau}$ is annual operating benefit excluding emissions damages, because $\bar{V}^{q^*_\tau}$ values fossil output at $p_f = 108.10 + \gamma_f \tau$; the marginal curve is the carbon price τ. The zero-emission point evaluates the same expression with the optimal mix without fossil backup from Table 7 in place of q^*_τ, for which emissions and the carbon-price term are zero.

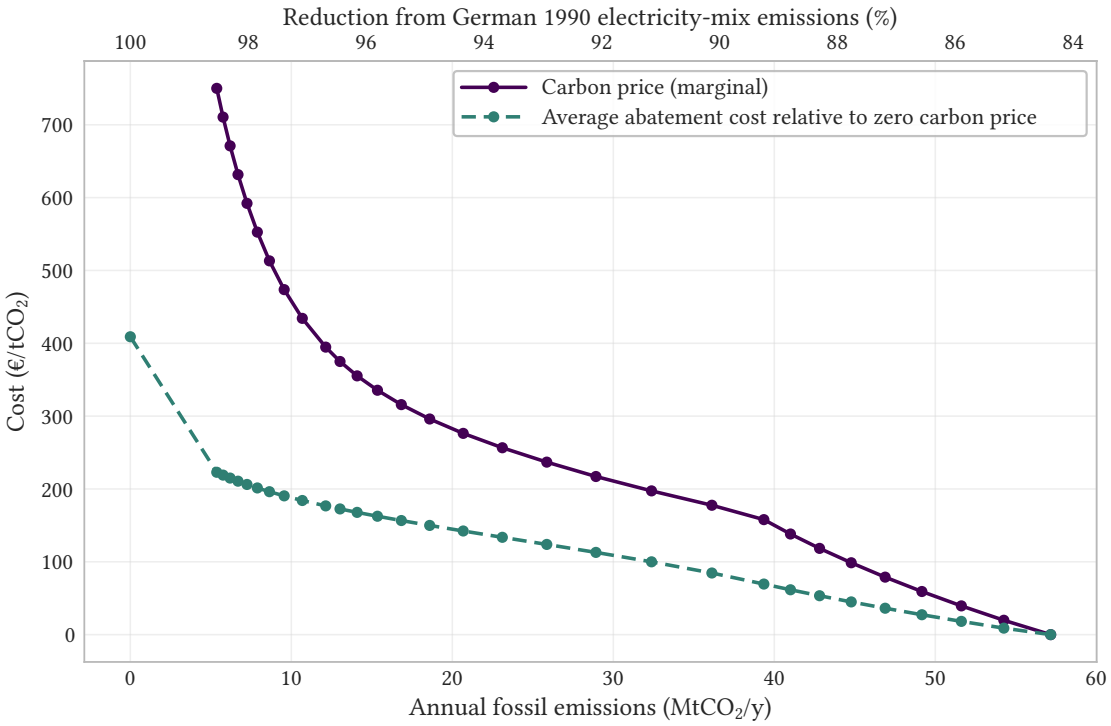


Figure 18: Marginal and average abatement costs. Each point is the optimal mix at one carbon price, which rises from right to left; the lower axis gives its annual fossil emissions and the upper axis their reduction from the 1990 emissions of the German electricity mix. The marginal abatement cost is the carbon price, and the average abatement cost is measured relative to the optimum at zero carbon price. The leftmost point is the zero-emission optimum without fossil backup.

The upper axis uses German electricity-mix emissions of 366 $MtCO_2/y$ in 1990 as its benchmark (Icha and Lauf, 2025). These decreased to 175 $MtCO_2/y$ in 2023 (Icha and Lauf, 2025). The optimal mix without a carbon price further reduces emissions to about 84% below the 1990 level. A further 10 percentage points are achievable with a carbon price of around 280 €/tCO_2. Completely phasing out natural gas requires extremely high carbon prices, even though the zero-emission mix (the leftmost point) features an average abatement cost of about 410 €/tCO_2. The optimum at 750 €/tCO_2 still emits 5.4 $MtCO_2/y$ at an average abatement cost of about 223 €/tCO_2, so the zero-emission point implies that eliminating these last tons costs about 2 200 €/tCO_2 on average. Early abatement replaces natural gas burned in many hours and is cheap per ton, whereas the last tons are emitted in rare drought hours, in which avoiding them requires storage capacity that is idle in most years (Section 7.2). The average curve therefore lies far below the marginal curve near zero emissions: it spreads the cost of the last tons over the cheap tons abated first. The welfare-adjusted levelized cost of the optimal mix with fossil backup, which values emissions at the respective carbon price, rises from 92.34 €/MWh without a carbon price

to 125.54 €/MWh at 750 €/tCO_2. Full decarbonization at 142.19 €/MWh (Table 8) therefore raises the welfare-adjusted levelized cost by about 50 €/MWh relative to the optimum without a carbon price and by about 17 €/MWh relative to the optimum at 750 €/tCO_2. Annual CO_2 emissions heavily depend on the weather once the carbon price is high. If biogas or imported carbon-neutral fuels could be used to cover shortfalls of the wind-solar-battery-hydrogen system, decarbonization would likely be substantially less costly and the marginal abatement cost would be bounded by the cost difference of producing the amount of electricity that emits one ton of CO_2 with such alternatives rather than with natural gas.

8 Discussion

I find that the optimal wind-solar-battery-hydrogen mix implies substantial increases in wind, solar, and hydrogen storage capacities while featuring very limited battery capacities, under perfect transmission and absent ramping constraints. Its welfare-adjusted levelized cost of 142.19 €/MWh at a 5% discount rate lies within the range, from 121 to 152 €/MWh, of the welfare-adjusted levelized costs of meeting the same demand with natural gas, hard coal, or lignite alone at recent fuel prices and the August 2026 carbon price (Appendix O). Both findings are largely in line with the literature on long-duration storage. With batteries as the only storage technology, the welfare-adjusted levelized cost doubles, to 284.61 €/MWh, because batteries are prohibitively costly for providing the large stockpile required. Without batteries but with hydrogen storage, it rises by only about 1%, to 143.84 €/MWh (Section 5.5), corroborating the conclusion that cheap long-duration storage is key for a 100% renewable electricity mix, even with 66% round-trip losses.

Introducing uncertainty as a continuous-time Markov process into an electricity dispatch and capacity choice model affects the optimal operation and the size of the hydrogen stockpile: under perfect foresight over the observed decade, the optimal mix holds markedly less hydrogen storage, while its renewable and conversion capacities differ less. Evaluating the optimal mix under uncertainty with perfect-foresight operation suggests that its annual operating value can differ by more than 8.5 B€/y across the 256 simulated paths, whereas the average value of perfect foresight is about 2.8 B€/y, contrasting with annualized system costs of more than 60 B€/y. This value is low because the optimal mix already holds a large hydrogen stockpile, which cheap cavern storage capacity makes affordable. For the 2016-2025 perfect-foresight optimum, which holds 23 rather than 35 TWh of hydrogen storage, the average value of perfect foresight is about 6.5 B€/y, more than twice as high (Section 6).

The shadow-price and consumption distributions are noticeably flatter in the fully renewable mix than in observed data, and the shadow price exceeds the price cap in about 23 hours per year. When the optimum is decentralized with price-taking firms, revenues discounted over 20 years are subject to substantial variation, suggesting that, even without a price cap, a capacity mechanism (Joskow and Tirole, 2007) may be beneficial to avoid large risk premia under spot pricing if investors require compensation for this revenue risk, which the risk-neutral model does not capture.

A higher price elasticity of demand substantially lowers hydrogen storage and conversion capacities and leaves battery capacity nearly unchanged: more flexible consumption reduces the need for a large reserve against prolonged low-production periods. For future research, this raises the question of separating risk aversion, that is, the aversion to volatile consumption implied by the curvature of the consumption benefit, from demand flexibility; this paper uses a single elasticity for both.

With a gas-turbine fossil backup, I find that expanding renewable capacities substantially is optimal, absent any carbon price. The optimal mix at zero carbon price reduces emissions by about 84% compared to 1990, with virtually no storage. Hydrogen storage enters the optimal mix at a carbon price between about 160 and 180 €/tCO_2. In the (presumably unlikely) case that even at a long horizon no carbon-neutral fuel can be imported, the marginal abatement cost increases sharply (above 750 €/tCO_2) as emissions approach zero. On the other hand, the average abatement cost of going from the optimal mix with fossil backup and without a carbon price to the optimal wind-solar-battery-hydrogen mix is about 410 €/tCO_2, and this step raises the welfare-adjusted levelized cost by about 50 €/MWh.

Battery capacity is by far the most cost-elastic capacity of the optimal mix: with low battery costs, it rises from 44 to 314 GWh (Section 5). Ramping constraints (on hydrogen charging, discharging, and the fossil backup) could give batteries a larger role in the optimal electricity mix; modeling them requires an additional state variable and is left for future research.

The results rest on further simplifications. Germany is modeled without electricity trade with neighboring countries.[38] All capacities are chosen at their full lifecycle cost, so existing assets, such as the pumped-storage capacity installed in 2025 (Table 19), do not enter the mix. Surplus renewable production is counted as consumption rather than curtailed (Section 2), which contributes to the flatter simulated consumption distribution. Wind production and costs are calibrated on onshore wind only. Reference consumption is derived from the observed 2025 load and is the same in every simulated year. Demand thus varies with the season but not with the weather within a season. If high demand tended to coincide with prolonged production shortfalls, for example during cold, calm spells, the model could understate scarcity and the required storage. Finally, the weather model somewhat understates the most extreme joint shortfalls of wind and solar production (Section 3) and the variation of wind production from

[38] Schlachtberger et al. (2017) find that, in a highly renewable European system, restricting transmission between countries raises total system cost by up to 30%, albeit with higher cost for hydrogen storage, which suggests that trade could lower the cost of a fully renewable German mix.

one winter to the next (Appendix E), which may lead the model to understate the hydrogen stockpile needed to cover a poor winter.

Acknowledgments

I thank my advisor, Christian Gollier, for all the feedback and support provided. Moreover, I thank Ingela Alger, Stefan Ambec, Charlotte de Cannière, Sylvain Chabé-Ferret, Fabrice Collard, Graham Macklin, Nour Meddahi, Anouch Missirian, Leticia Pieraerts, Mathias Reynaert, and Nicolas Treich for their guidance and useful comments. Finally, I thank all members of the TSE Environmental Group and participants at various conferences and seminars for their questions and comments.

Disclosure of Generative AI Use

I used generative AI tools for

- writing and debugging code, in particular adapting an initial model with two state variables (wind and the level of a single storage technology) to more state variables and seasonality, writing GPU code, switching to a semi-Lagrangian scheme, generating figures, and restructuring the code,
- discussing and developing ideas while extending the initial model,
- literature search,
- drafting and editing.

AI tools were used between 2023 and 2026, and the variety of tools used reflects the development of and learning about AI in this period. Most importantly, OpenAI Codex, Claude Code, GitHub Copilot, oh-my-pi, and OpenCode were used. The latter two were used with both OpenAI Codex models and various open-weight models, most importantly GLM 5.2, GLM 5.3 (-Flash), DeepSeek v4(.1) Flash, and Kimi K3.

I reviewed and verified all AI-assisted output, and I am solely responsible for the content of this paper.

References

Aichner, N. et al. (no date) *Energy Storage*. Open Energy Tracker. Available at: https://openenergytracker.org/en/docs/germany/storage/ (Accessed: August 15, 2026).

Ailliot, P. and Monbet, V. (2012) "Markov-switching autoregressive models for wind time series," *Environmental Modelling & Software*, 30. Available at: https://doi.org/10.1016/j.envsoft.2011.10.011.

Ambec, S. and Crampes, C. (2019) "Decarbonizing Electricity Generation with Intermittent Sources of Energy," *Journal of the Association of Environmental and Resource Economists*, 6(6). Available at: https://doi.org/10.1086/705536.

Andrés-Cerezo, D. and Fabra, N. (2023) "Storing power: market structure matters," *The RAND Journal of Economics*, 54(1). Available at: https://doi.org/10.1111/1756-2171.12429.

Antweiler, W. and Muesgens, F. (2025) "The new merit order: The viability of energy-only electricity markets with only intermittent renewable energy sources and grid-scale storage," *Energy Economics*, 145. Available at: https://doi.org/10.1016/j.eneco.2025.108439.

Asamov, T., Salas, D.F. and Powell, W.B. (2016) *SDDP vs. ADP: The Effect of Dimensionality in Multistage Stochastic Optimization for Grid Level Energy Storage*. arXiv. Available at: https://doi.org/10.48550/arXiv.1605.01521.

Badosa, J. et al. (2018) "Day-Ahead Probabilistic Forecast of Solar Irradiance: A Stochastic Differential Equation Approach," in P. Drobinski et al. (eds.) *Renewable Energy: Forecasting and Risk Management*. Cham: Springer International Publishing. Available at: https://doi.org/10.1007/978-3-319-99052-1_4.

Bermúdez, A. and Padín, I. (2024) "Stochastic control with state constraints via the Fokker–Planck equation. Application to renewable energy plants with batteries," *Comptes Rendus. Mécanique*, 351(S1). Available at: https://doi.org/10.5802/crmeca.236.

Blanchard, A. and Megy, C. (2026) "Battery and hydrogen storage: Complements or substitutes? A German case study," *Energy Economics*, 158. Available at: https://doi.org/10.1016/j.eneco.2026.109331.

Board of Governors of the Federal Reserve System (2026) *U.S. Dollars to Euro Spot Exchange Rate*. Available at: https://fred.stlouisfed.org/series/EXUSEU (Accessed: October 13, 2026).

Bobtcheff, C. (2011) "Optimal Dynamic Management of a Renewable Energy Source under Uncertainty," *Annals of Economics and Statistics*, (103/104). Available at: https://doi.org/10.2307/41615497.

Brown, T., Neumann, F. and Riepin, I. (2025) "Price Formation without Fuel Costs: The Interaction of Demand Elasticity with Storage Bidding," *Energy Economics*, 147. Available at: https://doi.org/10.1016/j.eneco.2025.108483.

Bundesnetzagentur (2025) *SMARD.de Benutzerhandbuch*. User manual. Bonn. Available at: https://www.smard.de/resource/blob/215828/3d407b3638ac86f6341aca9ca8f3d7ec/smard-benutzerhandbuch-02-2025-data.pdf (Accessed: October 13, 2026).

Byrd, R.H. et al. (1995) "A Limited Memory Algorithm for Bound Constrained Optimization," *SIAM Journal on Scientific Computing*, 16(5). Available at: https://doi.org/10.1137/0916069.

Caglayan, D.G. et al. (2020) "Technical potential of salt caverns for hydrogen storage in Europe," *International Journal of Hydrogen Energy*, 45(11). Available at: https://doi.org/10.1016/j.ijhydene.2019.12.161.

Cambridge Economic Policy Associates (2018) *Study on the Estimation of the Value of Lost Load of Electricity Supply in Europe*. Final Report ACER/OP/DIR/08/2013/LOT 2/RFS 10. Ljubljana. Available at: https://www.acer.europa.eu/en/Electricity/Infrastructure_and_network%20development/Infrastructure/Documents/CEPA%20study%20on%20the%20Value%20of%20Lost%20Load%20in%20the%20electricity%20supply.pdf.

Cao, X.-R. (2007) *Stochastic Learning and Optimization*. Boston, MA: Springer US. Available at: https://doi.org/10.1007/978-0-387-69082-7.

Carmona, R. and Ludkovski, M. (2010) "Valuation of energy storage: an optimal switching approach," *Quantitative Finance*, 10(4). Available at: https://doi.org/10.1080/14697680902946514.

Chen, Z. and Forsyth, P.A. (2008) "A Semi-Lagrangian Approach for Natural Gas Storage Valuation and Optimal Operation," *SIAM Journal on Scientific Computing*, 30(1). Available at: https://doi.org/10.1137/060672911.

Crampes, C. and Trochet, J.-M. (2019) "Economics of stationary electricity storage with various charge and discharge durations," *Journal of Energy Storage*, 24. Available at: https://doi.org/10.1016/j.est.2019.04.020.

Danish Energy Agency (2026a) *Data Sheet for Electricity and District Heat Production*. Available at: https://ens.dk/en/analyses-and-statistics/technology-data-generation-electricity-and-district-heating.

Danish Energy Agency (2026c) *Data Sheets for Energy Carrier Generation and Conversion*. Available at: https://ens.dk/en/analyses-and-statistics/technology-data-renewable-fuels.

Danish Energy Agency (2026b) *Datasheet for Energy Storage*. Available at: https://ens.dk/en/analyses-and-statistics/technology-data-energy-storage.

Deaton, A. (1991) "Saving and Liquidity Constraints," *Econometrica*, 59(5). Available at: http://www.jstor.org/stable/2938366 (Accessed: January 22, 2026).

Deutscher Bundestag (2019) *Bundes-Klimaschutzgesetz (KSG)*. Available at: https://www.gesetze-im-internet.de/ksg/.

Dowling, J.A. et al. (2020) "Role of Long-Duration Energy Storage in Variable Renewable Electricity Systems," *Joule*, 4(9). Available at: https://doi.org/10.1016/j.joule.2020.07.007.

Durmaz, T. (2016) *Precautionary Storage in Electricity Markets*. SSRN. Available at: https://doi.org/10.2139/ssrn.2738004.

European Network of Transmission System Operators for Electricity (ENTSO-E) (2026) *ENTSO-E Transparency Platform*. Available at: https://transparency.entsoe.eu/.

European Parliament (2021) *Regulation (EU) 2021/1119 of the European Parliament and of the Council of 30 June 2021 establishing the framework for achieving climate neutrality and amending Regulations (EC) No 401/2009 and (EU) 2018/1999 ('European Climate Law')*. Available at: http://data.europa.eu/eli/reg/2021/1119/oj (Accessed: October 27, 2026).

European Union Agency for the Cooperation of Energy Regulators (2026) *Decision No 02/2026 of the European Union Agency for the Cooperation of Energy Regulators of 4 February 2026 on the Nominated Electricity Market Operators' Proposal for the Amendment of the Harmonised Maximum and Minimum Clearing Price Methodology for Single Day-Ahead Coupling*. ACER Decision 02/2026. Ljubljana. Available at: https://www.acer.europa.eu/sites/default/files/documents/Individual%20Decisions_annex/ACER-Decision-02-2026-AnnexI.pdf.

Eurostat (2026a) *Harmonised index of consumer prices (HICP) - ECOICOP ver.2 - indices and rates of change, annual data*. Eurostat. Available at: https://ec.europa.eu/eurostat/databrowser/view/prc_hicp_ainr/default/table?lang=en (Accessed: October 21, 2026).

Eurostat (2026b) *Harmonised index of consumer prices (HICP) - ECOICOP ver.2 - indices and rates of change, monthly data*. Eurostat. Available at: https://ec.europa.eu/eurostat/databrowser/view/prc_hicp_minr/default/table?lang=en (Accessed: October 21, 2026).

Falcone, M. and Ferretti, R. (2013) *Semi-Lagrangian Approximation Schemes for Linear and Hamilton—Jacobi Equations*. Philadelphia, PA: Society for Industrial, Applied Mathematics. Available at: https://doi.org/10.1137/1.9781611973051.

Fleming, W.H. and Soner, H.M. (2006) *Controlled Markov processes and viscosity solutions*. 2. ed., [Nachdr.]. New York,NY]: Springer (Stochastic modelling and applied probability).

Fritsch, F.N. and Butland, J. (1984) "A Method for Constructing Local Monotone Piecewise Cubic Interpolants," *SIAM Journal on Scientific and Statistical Computing*, 5(2). Available at: https://doi.org/10.1137/0905021.

Gentile, C. (2024) "Relying on Intermittency: Clean Energy, Storage, and Innovation in a Macro Climate Model."

Geske, J. and Green, R. (2020) "Optimal Storage, Investment and Management under Uncertainty: It is Costly to Avoid Outages!," *The Energy Journal*, 41(2). Available at: https://doi.org/10.5547/01956574.41.2.jges.

Gowrisankaran, G., Reynolds, S.S. and Samano, M. (2016) "Intermittency and the Value of Renewable Energy," *Journal of Political Economy*, 124(4). Available at: https://doi.org/10.1086/686733.

Hamilton, J.D. (1989) "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle," *Econometrica*, 57(2). Available at: https://doi.org/10.2307/1912559.

Helm, C. and Mier, M. (2019) "On the efficient market diffusion of intermittent renewable energies," *Energy Economics*, 80. Available at: https://doi.org/10.1016/j.eneco.2019.01.017.

Helm, C. and Mier, M. (2021) "Steering the energy transition in a world of intermittent electricity supply: Optimal subsidies and taxes for renewables and storage," *Journal of Environmental Economics and Management*, 109. Available at: https://doi.org/10.1016/j.jeem.2021.102497.

Hirth, L. (2013) "The Market Value of Variable Renewables: The Effect of Solar Wind Power Variability on Their Relative Price," *Energy Economics*, 38. Available at: https://doi.org/10.1016/j.eneco.2013.02.004.

Hirth, L., Khanna, T.M. and Ruhnau, O. (2024) "How aggregate electricity demand responds to hourly wholesale price fluctuations," *Energy Economics*, 135. Available at: https://doi.org/10.1016/j.eneco.2024.107652.

Holland, S., Mansur, E. and Yates, A. (2022) *Decarbonization and Electrification in the Long Run*. technical report w30082. Cambridge, MA. Available at: https://doi.org/10.3386/w30082.

Howell, S.D. et al. (2009) "A partial differential equation system for modelling stochastic storage in physical systems with applications to wind power generation," *IMA Journal of Management Mathematics*, 22(3). Available at: https://doi.org/10.1093/imaman/dpp022.

Icha, P. and Lauf, T. (2025) *Entwicklung der spezifischen Treibhausgas-Emissionen des deutschen Strommix in den Jahren 1990 - 2024*. technical report. Available at: https://doi.org/10.60810/OPENUMWELT-7844.

International Energy Agency (2025) "Global Energy and Climate Model 2025."

International Monetary Fund (2026) *Global price of Natural gas, EU*. Available at: https://fred.stlouisfed.org/series/PNGASEUUSDM (Accessed: October 13, 2026).

Ishii, H. and Loreti, P. (2002) "A class of stochastic optimal control problems with state constraint," *Indiana University Mathematics Journal*.

Johnson, P., Howell, S. and Duck, P. (2017) "Partial differential equation methods for stochastic dynamic optimization: an application to wind power generation with energy storage," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 375(2100). Available at: https://doi.org/10.1098/rsta.2016.0301.

Jones, M.C. and Pewsey, A. (2009) "Sinh-arcsinh distributions," *Biometrika*, 96(4). Available at: https://doi.org/10.1093/biomet/asp053.

Joskow, P. and Tirole, J. (2007) "Reliability and Competitive Electricity Markets," *The RAND Journal of Economics*, 38(1). Available at: https://doi.org/10.1111/j.1756-2171.2007.tb00044.x.

Joskow, P.L. (2011) "Comparing the Costs of Intermittent and Dispatchable Electricity Generating Technologies," *American Economic Review*, 101(3). Available at: https://doi.org/10.1257/aer.101.3.238.

Juhrich, K. (2022) *CO2 Emission Factors for Fossil Fuels*. technical report 29/2022.

Junge, C., Mallapragada, D. and Schmalensee, R. (2022) "Energy Storage Investment and Operation in Efficient Electric Power Systems," *The Energy Journal*, 43(6). Available at: https://doi.org/10.5547/01956574.43.6.cjun.

Keske, C., Steffen, B. and Parsons, J.E. (2026) "How Uncertainty Shapes Electricity Storage Decisions: Dispatch Policies and Capacity Portfolios under Renewable Drought Risk." MIT Center for Energy, Environmental Policy Research.

Kittel, M. and Schill, W.-P. (2024) "Measuring the Dunkelflaute: how (not) to analyze variable renewable energy shortage," *Environmental Research: Energy*, 1(3). Available at: https://doi.org/10.1088/2753-3751/ad6dfc.

Kittel, M., Roth, A. and Schill, W.-P. (2026) "Long-Duration Electricity Storage Needs for Coping with Dunkelflaute Events in Europe," *Nature Communications*, 17(1). Available at: https://doi.org/10.1038/s41467-026-72681-5.

Kost, C. et al. (2024) *Stromgestehungskosten Erneuerbare Energien*. Studie. Freiburg. Available at: https://www.ise.fraunhofer.de/content/dam/ise/de/documents/publications/studies/DE2024_ISE_Studie_Stromgestehungskosten_Erneuerbare_Energien.pdf (Accessed: October 13, 2026).

Kuhn, H.W. (1960) "Some Combinatorial Lemmas in Topology," *IBM Journal of Research and Development*, 4(5). Available at: https://doi.org/10.1147/rd.45.0518.

Kushner, H.J. and Dupuis, P. (2001) *Numerical Methods for Stochastic Control Problems in Continuous Time*. Edited by I. Karatzas and M. Yor. New York, NY: Springer New York (Stochastic Modelling and Applied Probability). Available at: https://doi.org/10.1007/978-1-4613-0007-6.

Lai, G., Margot, F. and Secomandi, N. (2010) "An Approximate Dynamic Programming Approach to Benchmark Practice-Based Heuristics for Natural Gas Storage Valuation," *Operations Research*, 58(3). Available at: https://doi.org/10.1287/opre.1090.0768.

Löhndorf, N., Wozabal, D. and Minner, S. (2013) "Optimizing Trading Decisions for Hydro Storage Systems Using Approximate Dual Dynamic Programming," *Operations Research*, 61(4). Available at: https://doi.org/10.1287/opre.2013.1182.

Lütkehus, I., Salecker, H. and Adlunger, K. (2013) *Potenzial der Windenergie an Land*. technical report.

Nelsen, R.B. (2006) *An introduction to copulas*. 2nd ed. New York: Springer (Springer series in statistics).

Öberg, S., Odenberger, M. and Johnsson, F. (2022) "Exploring the competitiveness of hydrogen-fueled gas turbines in future energy systems," *International Journal of Hydrogen Energy*, 47(1). Available at: https://doi.org/10.1016/j.ijhydene.2021.10.035.

Picarelli, A. and Vargiolu, T. (2021) "Optimal management of pumped hydroelectric production with state constrained optimal control," *Journal of Economic Dynamics and Control*, 126. Available at: https://doi.org/10.1016/j.jedc.2020.103940.

Pommeret, A. and Schubert, K. (2022) "Optimal energy transition with variable and intermittent renewable electricity generation," *Journal of Economic Dynamics and Control*, 134. Available at: https://doi.org/10.1016/j.jedc.2021.104273.

Puterman, M.L. (1994) *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. 1st ed. Wiley (Wiley Series in Probability and Statistics). Available at: https://doi.org/10.1002/9780470316887.

Rennert, K. et al. (2022) "Comprehensive evidence implies a higher social cost of CO2," *Nature*, 610(7933). Available at: https://doi.org/10.1038/s41586-022-05224-9.

Ruhnau, O. and Qvist, S. (2022) “Storage Requirements in a 100% Renewable Electricity System: Extreme Events and Inter-Annual Variability,” *Environmental Research Letters*, 17(4). Available at: https://doi.org/10.1088/1748-9326/ac4dc8.

Salas, D.F. and Powell, W.B. (2018) “Benchmarking a Scalable Approximate Dynamic Programming Algorithm for Stochastic Control of Grid-Level Energy Storage,” *INFORMS Journal on Computing*, 30(1). Available at: https://doi.org/10.1287/ijoc.2017.0768.

Scheinkman, J.A. and Schechtman, J. (1983) “A Simple Competitive Model with Production and Storage,” *The Review of Economic Studies*, 50(3). Available at: https://doi.org/10.2307/2297674.

Schlachtberger, D.P. et al. (2017) “The Benefits of Cooperation in a Highly Renewable European Electricity Network,” *Energy*, 134. Available at: https://doi.org/10.1016/j.energy.2017.06.004.

Schmidt, F. (2025) *On long-duration storage, weather uncertainty and limited foresight*. arXiv. Available at: https://doi.org/10.48550/arXiv.2505.12538.

Secomandi, N. (2010) “Optimal Commodity Trading with a Capacitated Storage Asset,” *Management Science*, 56(3). Available at: https://doi.org/10.1287/mnsc.1090.1049.

Sepulveda, N.A. et al. (2021) “The Design Space for Long-Duration Energy Storage in Decarbonized Power Systems,” *Nature Energy*, 6(5). Available at: https://doi.org/10.1038/s41560-021-00796-8.

Sepulveda, N.A. et al. (2018) “The Role of Firm Low-Carbon Electricity Resources in Deep Decarbonization of Power Generation,” *Joule*, 2(11). Available at: https://doi.org/10.1016/j.joule.2018.08.006.

Shen, D., Ilic, M. and Parsons, J. (2026) *Dynamic Storage Operation Under Uncertainty and the Reliability Externality: Implications for Capacity Investments*. arXiv. Available at: https://doi.org/10.48550/ARXIV.2607.27021.

Sinn, H.-W. (2017) “Buffering volatility: A study on the limits of Germany's energy revolution,” *European Economic Review*, 99. Available at: https://doi.org/10.1016/j.euroecorev.2017.05.007.

Sisternes, F.J. de, Jenkins, J.D. and Botterud, A. (2016) “The Value of Energy Storage in Decarbonizing the Electricity Sector,” *Applied Energy*, 175. Available at: https://doi.org/10.1016/j.apenergy.2016.05.014.

Soner, H.M. (1986a) “Optimal Control with State-Space Constraint I,” *SIAM Journal on Control and Optimization*, 24(3). Available at: https://doi.org/10.1137/0324032.

Soner, H.M. (1986b) “Optimal Control with State-Space Constraint. II,” *SIAM Journal on Control and Optimization*, 24(6). Available at: https://doi.org/10.1137/0324067.

Stacy, E.W. (1962) “A Generalization of the Gamma Distribution,” *The Annals of Mathematical Statistics*, 33(3). Available at: https://doi.org/10.1214/aoms/1177704481.

Statistisches Bundesamt (no date) *Einfuhr von Steinkohle für das erste Halbjahr 2026*. Available at: https://www.destatis.de/DE/Themen/Branchen-Unternehmen/Energie/Verwendung/Tabellen/einfuhr-steinkohle-jaehrlich.html (Accessed: October 13, 2026).

Steffen, B. and Weber, C. (2013) "Efficient storage capacity in power systems with thermal and renewable generation," *Energy Economics*, 36. Available at: https://doi.org/10.1016/j.eneco.2012.11.007.

Ueckerdt, F. et al. (2013) "System LCOE: What are the costs of variable renewables?," *Energy*, 63. Available at: https://doi.org/10.1016/j.energy.2013.10.072.

Umweltbundesamt (2026) *Auctioning (EU-ETS 1): German Auctioning of Emission Allowances – August 2026*. technical report.

Varin, C., Reid, N. and Firth, D. (2011) "An overview of composite likelihood methods," *Statistica Sinica*.

Wang, Z., Wang, X. and Liu, W. (2023) "Genetic least square estimation approach to wind power curve modelling and wind power prediction," *Scientific Reports*, 13(1). Available at: https://doi.org/10.1038/s41598-023-36458-w.

Weber, P. and Woerman, M. (2024) "Intermittency or Uncertainty? Impacts of Renewable Energy in Electricity Markets," *Journal of the Association of Environmental and Resource Economists*, 11(6). Available at: https://doi.org/10.1086/729841.

Xi, X., Sioshansi, R. and Marano, V. (2014) "A stochastic dynamic programming model for co-optimization of distributed energy storage," *Energy Systems*, 5(3). Available at: https://doi.org/10.1007/s12667-013-0100-6.

Young, E.R. (2010) "Solving the incomplete markets model with aggregate uncertainty using the Krusell–Smith algorithm and non-stochastic simulations," *Journal of Economic Dynamics and Control*, 34(1). Available at: https://doi.org/10.1016/j.jedc.2008.11.010.

Zarate-Minano, R., Mele, F.M. and Milano, F. (2016) "SDE-based wind speed models with Weibull distribution and exponential autocorrelation," in *2016 IEEE Power and Energy Society General Meeting (PESGM)*. Boston, MA, USA: IEEE. Available at: https://doi.org/10.1109/PESGM.2016.7741754.

Zerrahn, A. and Schill, W.-P. (2017) "Long-run power storage requirements for high shares of renewables: review and a new model," *Renewable and Sustainable Energy Reviews*, 79. Available at: https://doi.org/10.1016/j.rser.2016.11.098.

Zerrahn, A., Schill, W.-P. and Kemfert, C. (2018) "On the Economics of Electrical Storage for Variable Renewable Energy Sources," *European Economic Review*, 108. Available at: https://doi.org/10.1016/j.euroecorev.2018.07.004.

Ziegler, M.S. et al. (2019) "Storage Requirements and Costs of Shaping Renewable Energy Toward Grid Decarbonization," *Joule*, 3(9). Available at: https://doi.org/10.1016/j.joule.2019.06.012.

A Electricity Consumption Benefits

This appendix states the monetary electricity consumption-benefit function described in Section 2 and calibrated in Section 3.

At phase ϑ, $U_\vartheta(c)$ is the monetary benefit from electricity consumption c, and $U'_\vartheta(c)$ is the marginal value of electricity. The formulation also admits a representative-consumer interpretation under quasilinear preferences: the consumer facing an electricity price p at phase ϑ solves

$$\max_{c\geq 0}\{U_\vartheta(c) - pc\}.$$

The reference consumption $c_{\text{ref}}(\vartheta)$ is the quantity chosen at the phase-independent reference price p_{ref}:

$$U'_\vartheta(c_{\text{ref}}(\vartheta)) = p_{\text{ref}}.$$

The price cap p_{cap} defines the curtailment threshold

$$\kappa = \left(\frac{p_{\text{ref}}}{p_{\text{cap}}}\right)^{|\varepsilon|}, \qquad k(\vartheta) = \kappa c_{\text{ref}}(\vartheta), \tag{29}$$

with the auxiliary coefficient

$$u_k(\vartheta) = \frac{p_{\text{ref}} c_{\text{ref}}(\vartheta)}{1+1/\varepsilon}\big(\kappa^{1+1/\varepsilon} - 1\big). \tag{30}$$

With the monetary consumption benefit normalized to zero at reference consumption, $U_\vartheta(c_{\text{ref}}(\vartheta)) = 0$, the consumption-benefit and marginal-value functions are

$$U_\vartheta(c) = \begin{cases} u_k(\vartheta) + p_{\text{cap}}(c - k(\vartheta)) - \frac{v_{\text{blackout}} - p_{\text{cap}}}{2k(\vartheta)}(c - k(\vartheta))^2, & c \leq k(\vartheta), \\ \frac{p_{\text{ref}} c_{\text{ref}}(\vartheta)^{-1/\varepsilon}}{1+1/\varepsilon}\big(c^{1+1/\varepsilon} - c_{\text{ref}}(\vartheta)^{1+1/\varepsilon}\big), & c > k(\vartheta), \end{cases} \tag{31}$$

$$U'_\vartheta(c) = \begin{cases} v_{\text{blackout}} - \frac{v_{\text{blackout}} - p_{\text{cap}}}{k(\vartheta)} c, & c \leq k(\vartheta), \\ p_{\text{ref}}\left(\frac{c}{c_{\text{ref}}(\vartheta)}\right)^{1/\varepsilon}, & c > k(\vartheta). \end{cases} \tag{32}$$

For $c > k(\vartheta)$, the marginal value of electricity follows inverse demand; for $c \leq k(\vartheta)$, it represents the marginal cost of unserved demand.

B Theoretical Details and HJB Characterization

In the absence of the drought regime, the Ornstein-Uhlenbeck process in (2) of Section 2 would have the identity matrix as its stationary covariance.

The drift and diffusion matrices therefore satisfy the stationary Lyapunov equation

$$\Psi_X + \Psi_X^T + QQ^T = 0. \tag{33}$$

For twice differentiable v, the generator of the weather process at regime r and phase ϑ is $\mathcal{L}_X v = (\Psi_X(x + rm(\vartheta))) \cdot \nabla_x v + \frac{1}{2}\,\mathrm{tr}(QQ^T D_x^2 v)$, where $D_x^2 v$ is the Hessian of v in the latent state.

The drought regime $R_t \in \{0, 1\}$ is a continuous-time Markov chain, independent of the Wiener process, with rates $\gamma_{01}, \gamma_{10} > 0$. For a function v of the regime, its generator is

$$\mathcal{J} v(r) = \gamma_{r,1-r}(v(1-r) - v(r)). \tag{34}$$

For a function of the full state, $\mathcal{J}$ acts on the regime at fixed (x, s, ϑ).

B.1 Relative Value

The relative value W^q of (10) in Section 2 is the expected net operating benefit in excess of the long-run hourly payoff ρ^q that accumulates from state ξ under an optimal stationary policy:

$$W^q(\xi) = \lim_{N\to\infty} \mathbb{E}_\xi \left[\int_0^{NT_Y - \vartheta} \left(U_{\theta_t}(c_t) - p_f f_t - \rho^q \right) dt \right]. \tag{35}$$

Subtracting ρ^q removes the long-run growth of the cumulative benefit, so the limit, which I assume exists, measures the advantage of starting from ξ. The upper limit $NT_Y - \vartheta$ lets every horizon end on January 1. Expected benefits are above ρ^q in some seasons and below it in others, so the accumulated excess depends on the season in which the horizon ends; fixing the end date removes this dependence.

B.2 State-Constrained HJB

The storage constraints make the operating problem a state-constrained stochastic control problem. Picarelli and Vargiolu (2021), in their controllable case, consider a closely related problem. The admissible action sets alone enforce the storage bounds; this relies on the inward-pointing condition (Soner, 1986a; 1986b).

The relative value W^q may lack the differentiability required for a classical solution of (10) in Section 2. In that case, the state-constrained viscosity solution framework provides a rigorous interpretation of the HJB; see Ishii and Loreti (2002). Because Brownian shocks act only on weather, the diffusion is degenerate in the storage and phase directions, which motivates the viscosity formulation (Picarelli and Vargiolu, 2021).

For $i \in \{H, B\}$, storage state constraints are represented by

$$A_i(s_i) = \begin{cases} [0, K_{i+}], & s_i = 0, \\ [-K_{i-}, K_{i+}], & 0 < s_i < \bar{S}_i, \\ [-K_{i-}, 0], & s_i = \bar{S}_i, \end{cases} \tag{36}$$

with $K_{i-} \geq 0$ and $K_{i+} \geq 0$.

The instantaneous net operating benefit is

$$\ell^q(\xi, \alpha, \nu) = U_\vartheta \left(y(\xi) + \nu - \sum_i \alpha_i \right) - p_f \nu.$$

The admissible action set is

$$A(\xi) = \left\{ (\alpha, \nu) : \alpha_i \in A_i(s_i), 0 \leq \nu \leq F, y(\xi) + \nu - \sum_i \alpha_i \geq 0 \right\}. \tag{37}$$

In the fossil-backup specification of Section 7, natural gas and hydrogen share the CCGT capacity, and the fossil bound in (37) becomes $0 \leq \nu \leq F + K_{H-} + \alpha_{H-}$. Fossil dispatch then lies in $[0, F + K_{H-}]$.

For sufficiently smooth W^q, the dynamic programming principle and Itô's formula yield the ergodic HJB (Fleming and Soner, 2006): the phase and storage dynamics contribute first derivatives, the Brownian weather process contributes $\mathcal{L}_X W^q$, and the regime contributes the jump term $\mathcal{J} W^q$.

For an action (α, ν) and a sufficiently smooth function v, the controlled generator is

$$\mathcal{L}^{\alpha,\nu} v = \partial_\vartheta v + \mathcal{L}_X v + \mathcal{J} v + \sum_i b_i(\alpha_i) \partial_{s_i} v, \tag{38}$$

where b_i is the storage-energy flow of Section 2. The ergodic HJB (10) in Section 2 then reads

$$\rho^q = \sup_{(\alpha,\nu) \in A(\xi)} \{ \ell^q(\xi, \alpha, \nu) + \mathcal{L}^{\alpha,\nu} W^q(\xi) \}. \tag{39}$$

The relative value function is periodic in the phase. No boundary condition is prescribed at the storage boundaries.

B.3 Verification

A standard verification argument (Fleming and Soner, 2006) links the HJB to the operating problem (7) in Section 2. Suppose that a scalar ρ, a phase-periodic function W, and a measurable stationary policy $\pi(\xi) \in A(\xi)$ satisfy (39) with ρ and W in place of ρ^q and W^q, that π attains the supremum at every state, and that the state equations under π have a solution in the state space $\mathcal{X}^q = \mathbb{R}^2 \times \{0, 1\} \times \left[0, \bar{S}_H\right] \times \left[0, \bar{S}_B\right] \times [0, T_Y)$ from every initial state. Suppose further that W is continuously differentiable in the phase and the stocks and twice continuously differentiable in the latent state, up to the storage boundaries; that W, its phase and stock derivatives,

and its second derivatives in the latent state grow at most quadratically in x, uniformly in the regime, the stocks, and the phase; and that its gradient in the latent state grows at most linearly. For every admissible control, Itô's formula and the HJB then give

$$\mathbb{E}_{\xi}\left[\int_{0}^{T} \ell^{q}(\Xi_{t}, a_{t}, f_{t})dt\right] \leq T\rho + W(\xi) - \mathbb{E}_{\xi}[W(\Xi_{T})], \tag{40}$$

with equality under π. The left side is finite because the net operating benefit is bounded: U_{ϑ} is bounded above and finite at zero consumption, and fossil dispatch is bounded. Weather is not controlled, every eigenvalue of Ψ_{X} has a negative real part, and the drought mean is bounded, so $\sup_{t\geq 0}\mathbb{E}_{\xi}\left[\|X_{t}\|^{2}\right] \leq M(1+\|x\|^{2})$ under every policy for some constant M. After division by T, the terminal term therefore vanishes as $T \to \infty$. Hence $\rho = \rho^{q}$ from every initial state, and π attains this payoff.

C Marginal Capacity Values through the HJB

The marginal value of a capacity is the change it makes to the value of optimally operating the entire system. This appendix characterizes that change for renewable capacities and for storage energy and power capacities, using the ergodic operating problem and capacity criterion (13) from Section 2. Here q is the full capacity vector in Section 2, including wind, solar, both storage technologies, their conversion power, and backup. Weather, consumption benefits, and efficiencies are held fixed. The objects of interest are the derivatives $\partial_{q_j}\rho^q$ with respect to the renewable capacities K_W and K_S and the storage capacities $\bar{S}_i$, K_{i+}, and K_{i-} for $i \in \{H, B\}$. The capacity optimization in Appendix F uses the discrete counterpart of these derivatives as the gradient of the NPV.

C.1 Statewise Marginal Capacity Values in the Operating HJB

The storage energy capacity $\bar{S}_i$ sets the range of the storage state s_i. To differentiate with respect to it, it is convenient to work with the normalized fullness and to define the value on the fixed unit interval,

$$u_i = \frac{s_i}{\bar{S}_i}, \quad \widehat{W}^q(x, r, u, \vartheta) = W^q\big(x, r, \big(\bar{S}_H u_H, \bar{S}_B u_B\big), \vartheta\big), \quad i \in \{H, B\}. \tag{41}$$

Here b_i is the storage-energy flow of (10) in Section 2. With $(N_W, N_S) = n_\vartheta(x)$ the weather state at (x, ϑ), output per GW is

$$\chi_W = \psi_W(\varsigma_W(\vartheta)N_W, \vartheta), \quad \chi_S = \psi_S(\varsigma_S(\vartheta)N_S, \vartheta)\phi_S(\vartheta).$$

The normalized ergodic HJB is

$$\rho^q = \partial_\vartheta \widehat{W}^q + \mathcal{L}_X \widehat{W}^q + \mathcal{J}\widehat{W}^q + \sup_{(c,\alpha,\nu)\in\hat{A}^q} \left\{ U_\vartheta(c) - p_f\nu + \sum_{i\in\{H,B\}} \frac{b_i(\alpha_i)}{\bar{S}_i}\partial_{u_i}\widehat{W}^q \right\}. \tag{42}$$

Here ∂_ϑ is the derivative with respect to the phase, $\mathcal{L}_X$ the weather generator, and $\mathcal{J}$ the regime generator from Appendix B. The feasible set $\hat{A}^q$ imposes $c + \alpha_H + \alpha_B = K_W\chi_W + K_S\chi_S + \nu$, nonnegative consumption, and the original power and storage constraints. In particular, $\alpha_i \geq 0$ at $u_i = 0$ and $\alpha_i \leq 0$ at $u_i = 1$.

Evaluate the following quantities under the optimal stationary policy:

$$\beta_i = \frac{b_i(a_i^q)}{\bar{S}_i}, \quad \lambda_i = \frac{\partial_{u_i}\widehat{W}^q}{\bar{S}_i} = \partial_{s_i}W^q, \quad \mathcal{L}^q v = \partial_\vartheta v + \mathcal{L}_X v + \mathcal{J}v + \sum_{i\in\{H,B\}} \beta_i \partial_{u_i} v.$$

Here λ_i is the shadow value of stored energy and β_i is the storage flow rate in normalized units, so $\mathcal{L}^q$ is the controlled generator of Appendix B at the optimal actions, written in normalized stocks. Optimal consumption is strictly positive when $K_W > 0$. Wind output is at

least $\underline{\psi}_W K_W > 0$, where $\underline{\psi}_W$ is the lower bound of the wind map ψ_W in Appendix E, and a unit of stored energy replaces at most $1/\omega_{i-}$ units of later consumption, each worth at most $U'_\vartheta(0) = v_{\text{blackout}}$. Because storage has round-trip losses, $\omega_{i+} < \omega_{i-}$, charging is worth less than v_{blackout}, so charging all output at zero consumption is never optimal. At the normalized operating state $\hat{\xi} = (x, r, u, \vartheta)$, the multiplier on the energy balance is therefore

$$p\big(\hat{\xi}\big) = U'_\vartheta\big(c^q\big(\hat{\xi}\big)\big).$$

The shadow price p, like λ_i and β_i, is a function of the operating state; state arguments and dependence on the fixed capacity vector q are suppressed below. Let $\hat{m}_j$ denote the direct capacity derivative of the optimized Hamiltonian, holding the value-function derivatives fixed, that is, the statewise marginal value of capacity q_j. For the equipment capacities,

$$\begin{aligned} &\hat{m}_{K_W} = p\chi_W, \quad \hat{m}_{K_S} = p\chi_S, \\ &\hat{m}_{K_{i+}} = \Lambda_{i+}, \quad \hat{m}_{K_{i-}} = \Lambda_{i-}, \quad \hat{m}_F = \Lambda_F, \\ &\hat{m}_{\bar{S}_i} = -\frac{b_i(a_i^q)\lambda_i}{\bar{S}_i}. \end{aligned} \tag{43}$$

The nonnegative multipliers Λ_{i+}, Λ_{i-}, and Λ_F value the charging, discharging, and fossil output limits and vanish when their respective limits are slack; at positive independent fossil capacity, $\Lambda_F = \big[p - p_f\big]_+$.

The statewise marginal values of storage power are scarcity rents. When a limit binds, an extra unit of power capacity is worth the marginal operating cost it relieves: $\omega_{i+}\lambda_i - p$ for charging and $p - \omega_{i-}\lambda_i$ for discharging.

The statewise marginal value of energy capacity is the net energy flow out of the storage, weighted by the shadow value of stored energy, per unit of capacity. It takes this form because the ergodic HJB depends on $\bar{S}_i$ only through the normalized flow $\beta_i = b_i(a_i^q)/\bar{S}_i$. A larger storage capacity slows filling while the storage is charged and slows emptying while it is discharged, so this marginal value is negative during charging and positive during discharging.

Along the selected optimal stationary policy for q, write $\hat{\Xi}_t = (X_t, R_t, u_t, \theta_t)$. The formulas below assume that ρ^q and $\widehat{W}^q$ are differentiable in the capacities and that each derivative $G_j = \partial_{q_j}\widehat{W}^q$ satisfies the regularity and growth conditions of the verification argument in Appendix B.

C.2 Marginal Values from the Time-Average Envelope

For a state function h, the averaging operator from (16) in Section 2 is

$$\mathcal{M}^q[h] = \lim_{T\to\infty} \frac{1}{T}\mathbb{E}_\xi\left[\int_0^T h\big(\hat{\Xi}_t\big)dt\right].$$

Products such as $p\chi_j$ are evaluated at the same state before averaging; in storage averages, b_i denotes the optimal net flow $b_i(a_i^q)$ as a function of that state. Differentiating the HJB along $\hat{\Xi}_t$, integrating over $[0,T]$, and dividing by T gives

$$\partial_{q_j}\rho^q = \frac{1}{T}\mathbb{E}_\xi\left[\int_0^T \left(\hat{m}_j\big(\hat{\Xi}_t\big) + \mathcal{L}^q G_j\big(\hat{\Xi}_t\big)\right)dt\right],$$

$$\frac{1}{T}\mathbb{E}_\xi\left[\int_0^T \mathcal{L}^q G_j\big(\hat{\Xi}_t\big)dt\right] = \frac{\mathbb{E}_\xi\left[G_j\big(\hat{\Xi}_T\big)\right] - \mathbb{E}_\xi\left[G_j\big(\hat{\Xi}_0\big)\right]}{T} \to 0, \tag{44}$$

$$\partial_{q_j}\rho^q = \mathcal{M}^q\left[\hat{m}_j\right].$$

Dynkin's formula gives the second line. The limit follows because the growth condition on G_j keeps $\mathbb{E}_\xi\left[G_j\big(\hat{\Xi}_T\big)\right]$ bounded in T. The marginal value of a capacity is therefore the long-run expected time average of its statewise marginal value. At an interior optimum, the investment condition (14) in Section 2 becomes $T_Y\mathcal{M}^{q^*}\left[\hat{m}_j\right] = \delta C_j$. For renewables, the marginal operating benefit is the expected long-run average of output valued at the state-dependent shadow price:

$$\begin{aligned}\partial_{K_j}\rho^q &= \mathcal{M}^q\left[p\chi_j\right] \\ &= \mathcal{M}^q[p]\mathcal{M}^q\left[\chi_j\right] + \mathrm{Cov}^q\big(p,\chi_j\big), \quad j \in \{W,S\},\end{aligned} \tag{45}$$

where $\mathrm{Cov}^q\big(p,\chi_j\big) = \mathcal{M}^q\left[p\chi_j\right] - \mathcal{M}^q[p]\mathcal{M}^q\left[\chi_j\right]$. Renewables are valued by production weighted by scarcity, so average output and its timing both matter. For storage, the same average gives

$$\partial_{K_{i+}}\rho^q = \mathcal{M}^q[\Lambda_{i+}], \quad \partial_{K_{i-}}\rho^q = \mathcal{M}^q[\Lambda_{i-}], \quad \partial_{\bar{S}_i}\rho^q = -\frac{\mathcal{M}^q[b_i\lambda_i]}{\bar{S}_i}. \tag{46}$$

D Lifecycle Costs, Prices, and Revenues

This appendix gives the lifecycle-cost calculation and the definitions of revenues, average sale prices, and certainty-equivalent consumption used in Section 2. It uses the policy functions c^q, f^q, and a_i^q, the shadow price p^q, and the long-run averages $\mathcal{M}^q$ of (16) defined there.

D.1 Lifecycle Capacity Costs

For component q_j of the capacity vector, let I_j be unit CAPEX, o_j annual fixed O&M, T_j the lifetime in years, and ζ_j the refurbishment fraction, that is, the fraction of the initial CAPEX paid again at the end of each lifetime in order to renew the installation. Then

$$\begin{aligned} C(q) &= \sum_j C_j q_j, \text{ where} \\ C_j &= I_j \left(1 + \zeta_j \sum_{k=1}^{\infty} e^{-kT_j\delta} \right) + \frac{o_j}{\delta} \\ &= I_j \left(1 + \zeta_j \frac{e^{-T_j\delta}}{1 - e^{-T_j\delta}} \right) + \frac{o_j}{\delta}. \end{aligned} \tag{47}$$

D.2 Consumption, Prices, and Revenues

For the flows considered in the results, I assume that the long-run averages (16) exist, are finite, and are independent of initial weather, regime, stocks, and the starting date. Write $c_t^q = c^q(\Xi_t)$, $p_t^q = p^q(\Xi_t)$, and $f_t^q = f^q(\Xi_t)$ for consumption, price, and fossil dispatch along the simulated operating paths. The long-run standard deviation of consumption is $\mathrm{SD}(c^q) = \left(\mathcal{M}^q \left[(c^q)^2 \right] - (\bar{c}^q)^2 \right)^{1/2}$.

Write $y_j^q(\xi)$ for the output in GW of renewable technology $j \in \{W, S\}$. Annual renewable revenues and storage revenues in k€/y are

$$\begin{aligned} \mathrm{Rev}_j^q &= T_Y \mathcal{M}^q \left[p^q y_j^q \right], \quad j \in \{W, S\}, \\ \mathrm{Rev}_i^q &= -T_Y \mathcal{M}^q [p^q a_i^q], \quad i \in \{H, B\}. \end{aligned} \tag{48}$$

Storage revenue is net of charging expenditure because $a_i^q > 0$ denotes charging. Dividing these revenues by the quantities sold gives quantity-weighted average prices in €/MWh. For wind and solar, the quantity is output; for storage, it is discharge, $-a_{i-}^q$, with the discharging component of a_i^q as in Section 2, so that the average price of storage is net of charging expenditure:

$$\begin{aligned} \bar{p}_j^q &= \frac{\mathcal{M}^q\left[p^q y_j^q\right]}{\mathcal{M}^q\left[y_j^q\right]}, \quad j \in \{W, S\}, \\ \bar{p}_i^q &= \frac{-\mathcal{M}^q\left[p^q a_i^q\right]}{\mathcal{M}^q\left[-a_{i-}^q\right]}, \quad i \in \{H, B\}. \end{aligned} \tag{49}$$

For hydrogen, cavern storage, electrolyzers and hydrogen CCGTs constitute a single system.

Certainty-equivalent consumption scales the reference profile. Let μ^q be the factor at which the deterministic profile $\mu^q c_{\mathrm{ref}}(\vartheta)$ yields the long-run expected benefit of consumption:

$$\frac{1}{T_Y}\int_0^{T_Y} U_\vartheta(\mu^q c_{\mathrm{ref}}(\vartheta))d\vartheta = \mathcal{M}^q\left[U_{\theta_t}(c_t^q)\right]. \tag{50}$$

Strict monotonicity of U_ϑ makes the left side strictly increasing in μ^q, so the factor is unique. Certainty-equivalent consumption is

$$\bar{c}_{\mathrm{CE}}^q = \mu^q \bar{c}_{\mathrm{ref}}. \tag{51}$$

E Calibration Details

This appendix gives the functional forms and estimation steps for renewable production and reference consumption. It compares the fitted marginal densities, intraday wind variation, and multi-week production with the data and shows hourly reference consumption on selected days.

E.1 Wind and Solar Marginals and Latent Dynamics

Let $X_t = (X_{Wt}, X_{St})$ denote the latent Ornstein-Uhlenbeck state in (2) of Section 2. Harmonic functions of the phase use the annual angle $2\pi\vartheta/T_Y$, with $\vartheta = 0$ at January 1, 00:00 UTC. The parameters of the weather model are specific to this appendix: symbols such as κ_S, δ_W, ε_W, and c_W are unrelated to κ, δ, ε, and consumption c in the main text.

Weather state and dependence. At phase ϑ, the latent state $x = (x_W, x_S)$ is mapped into the weather state $n_\vartheta(x)$ of Section 2 by

$$n_\vartheta(x) = \big(x_W, \Phi^{-1}\big(H_\vartheta^{-1}(\Phi(x_S) \mid \Phi(x_W))\big)\big), \quad H_\vartheta(v \mid u) = \partial_u C_\vartheta(u, v), \tag{52}$$

where Φ is the standard-normal distribution function and $H_\vartheta(\cdot \mid u)$ is the conditional distribution function implied by the copula C_ϑ when its first argument equals u. The wind component equals that of the latent state, and the solar component is placed at the same quantile of its conditional distribution given the wind component as the latent solar component in its standard-normal distribution. In the absence of the drought regime, $\Phi(X_{Wt})$ and $\Phi(X_{St})$ are independent and uniformly distributed, so both components of the weather state are standard normal, with joint distribution

$$\mathbb{P}(N_{Wt} \le a, N_{St} \le b) = C_{\theta_t}(\Phi(a), \Phi(b)). \tag{53}$$

The copula is a mixture with seasonal weight $w_C(\vartheta)$,

$$C_\vartheta = (1 - w_C(\vartheta))C^{\mathrm{Ga}}_{\zeta(\vartheta)} + w_C(\vartheta)C^{\mathrm{Cl}}_{\gamma_C}, \tag{54}$$

of the Gaussian copula $C^{\mathrm{Ga}}_\zeta(u, v) = \Phi_\zeta(\Phi^{-1}(u), \Phi^{-1}(v))$, where Φ_ζ is the bivariate standard-normal distribution function with correlation ζ, and the Clayton copula $C^{\mathrm{Cl}}_{\gamma_C}(u, v) = (u^{-\gamma_C} + v^{-\gamma_C} - 1)^{-1/\gamma_C}$ with parameter $\gamma_C > 0$ (Nelsen, 2006). The correlation $\zeta(\vartheta)$ is the hyperbolic tangent and the weight $w_C(\vartheta)$ the logistic function of a first-order harmonic. With $w_C = 0$ and in the absence of the drought regime, the weather state is bivariate standard normal with correlation $\zeta(\vartheta)$.

Drought regime. The drought shift is

$$D(\vartheta) = \frac{\bar{D}}{1 + \exp\!\left(-\left(a_D + b_D^{s} \sin\!\left(\tfrac{2\pi\vartheta}{T_Y}\right) + b_D^{c} \cos\!\left(\tfrac{2\pi\vartheta}{T_Y}\right)\right)\right)},$$
$$m(\vartheta) = D(\vartheta)\left(1, \frac{1-\zeta(\vartheta)}{\sqrt{1-\zeta(\vartheta)^2}}\right). \tag{55}$$

Under the Gaussian copula, the solar component is $N_S = \zeta X_W + \sqrt{1-\zeta^2} X_S$, so at the drought mean $X = -m(\vartheta)$ both components equal $-D(\vartheta)$. The seasonal weight $D/\bar{D}$ is about 0.8 in winter and 0.15 in summer, so the regime acts mainly in winter. At the drought mean and before scaling, the wind component of the weather state is lowered by up to 1.01 standard deviations in winter[39] but both components only by 0.18 in summer. Because the solar scale is small in winter, the scaled solar component is lowered most in spring and autumn. After the regime switches, production adjusts at the speed at which ordinary weather fluctuations decay, with decay times of 26 and 34 hours.[40]

Weather-state scales. The solar scale $\varsigma_S(\vartheta)$ of (3) in Section 2 is the exponential of a second-order harmonic of the annual phase. The wind scale $\varsigma_W(\vartheta)$ is the exponential of the sum of a first-order harmonic and the daily cycle $d_\sigma(\vartheta)$ defined below.

The wind capacity-factor map, motivated in Section 3, is

$$\psi_W(z,\vartheta) = \underline{\psi}_W + \frac{\bar{\psi}_W - \underline{\psi}_W}{1 + \exp(-\gamma_W(\sinh(\delta_W \operatorname{arsinh} z - \varepsilon_W) + \mu_W(\vartheta) - c_W))}, \tag{56}$$

where

$$\mu_W(\vartheta) = \mu_W^{s} \sin\!\left(\frac{2\pi\vartheta}{T_Y}\right) + \mu_W^{c} \cos\!\left(\frac{2\pi\vartheta}{T_Y}\right) + d_\mu(\vartheta). \tag{57}$$

The sinh-arcsinh transformation (Jones and Pewsey, 2009) with skewness ε_W and tail weight δ_W reduces to the identity for $\varepsilon_W = 0$ and $\delta_W = 1$. The bounds $\underline{\psi}_W$ and $\bar{\psi}_W$ are set from the observed range of hourly capacity factors rather than estimated: $\underline{\psi}_W$ is half the observed minimum, and $\bar{\psi}_W$ lies halfway between the observed maximum and 1, so both lie slightly outside the observed range. The daily cycles of the wind map's argument and of its scale have the same form,

$$d_l(\vartheta) = \left(1 + e_l^{s} \sin\!\left(\frac{2\pi\vartheta}{T_Y}\right) + e_l^{c} \cos\!\left(\frac{2\pi\vartheta}{T_Y}\right)\right)$$
$$\times \sum_{j=1}^{2}\left(g_{lj}^{s} \sin\!\left(\frac{2\pi j\vartheta}{24}\right) + g_{lj}^{c} \cos\!\left(\frac{2\pi j\vartheta}{24}\right)\right), \quad l \in \{\mu, \sigma\}, \tag{58}$$

[39] The solar component, under the fitted copula, is lowered by about 1.2-1.3 standard deviations.

[40] A decay time is the time in which a deviation falls to $1/e$ of its initial size; there is one for each eigenvalue of the drift matrix Ψ_X.

where ϑ modulo 24 is the UTC time of day. The two harmonics describe the daily cycle, and the annual modulation lets its amplitude vary over the year. Observed calm hours cluster in the morning and end when the wind rises during the day; without a daily cycle, the model can match the frequency of calm hours only with too many entirely calm days. The cycle of the scale is needed because the daily cycle is concentrated in the lower half of the distribution: in summer, the 5% and 25% quantiles fall by about half in the morning, whereas the 95% quantile barely moves.

The solar map, which gives the hourly solar capacity factor divided by the intraday solar profile, is a generalized Gamma quantile (Stacy, 1962) with a seasonal shape parameter and a seasonal exponent:

$$\begin{aligned} \psi_S(z,\vartheta) &= \sigma_\Gamma \kappa_S(\vartheta) \left(\frac{F^{-1}_{\Gamma(\kappa_S(\vartheta),1)}(\Phi(z))}{\kappa_S(\vartheta)} \right)^{1/\iota_S(\vartheta)}, \\ \log \iota_S(\vartheta) &= \iota_0 + \iota^s \sin\left(\frac{2\pi\vartheta}{T_Y}\right) + \iota^c \cos\left(\frac{2\pi\vartheta}{T_Y}\right). \end{aligned} \tag{59}$$

For $\iota_S = 1$, the map is the Gamma quantile with shape $\kappa_S(\vartheta)$ and scale σ_Γ. In summer, the estimated ι_S is about 8, and the power $1/\iota_S$ compresses the distribution and skews it to the left, so that most hours lie somewhat below a clear-sky ceiling. The seasonal shape parameter is

$$\kappa_S(\vartheta) = \kappa_0 + \kappa^s \sin\left(\frac{2\pi\vartheta}{T_Y}\right) + \kappa^c \cos\left(\frac{2\pi\vartheta}{T_Y}\right) + \kappa^s_2 \sin\left(\frac{4\pi\vartheta}{T_Y}\right) + \kappa^c_2 \cos\left(\frac{4\pi\vartheta}{T_Y}\right). \tag{60}$$

Estimation. All parameters of the maps, the copula, the latent dynamics, and the drought regime other than the wind-map bounds $\underline{\psi}_W$ and $\bar{\psi}_W$, 50 in total, are estimated jointly from hourly German capacity factors for 2015-2025 by composite maximum likelihood. The composite likelihood sums the log densities of the later observation given the earlier one in pairs of hourly observations 24, 72, and 168 hours apart. These lags emphasize persistence over days and weeks, which matters most for storage. For each pair, the density mixes over the regime at both dates, using the regime transition probability over the lag. Given the regimes at both dates, the latent state at the later date is treated as normal: its conditional mean includes the expected pull toward the drought mean along the regime path between the two dates, given its endpoints, and the additional variance due to the uncertain regime path is neglected. At the earlier date, the latent state is treated as normal around the target of its regime, 0 or $-m(\vartheta)$, with identity covariance. Combined with the stationary regime probabilities, these densities give the probabilities of each regime at the earlier date given the earlier observation; the densities themselves do not enter the composite likelihood. Solar observations enter only in hours in which the intraday solar profile

is at least half of its daily mean; in other hours, the pair density uses wind alone. Pairs with solar observations are weighted so that every day carries the same weight regardless of its number of daylight hours, and pairs that end in an hour in which both wind and solar fall below the 10% quantiles of their calendar month receive twice this weight. The doubled weight is a calibration choice rather than part of the likelihood: without it, the estimate understates the frequency of joint winter droughts. The fit of the share of winter days on which both daily capacity factors fall below their 10% quantiles (Section 3) is therefore partly targeted by the estimation rather than an independent check, and the sensitivity of the optimal storage capacities to this weight is not quantified. Because the composite likelihood treats overlapping pairs as independent, no standard errors are reported. Table 20 reports the estimates.

Component	Parameter	Estimate
Latent drift	Ψ_X (per hour)	$\begin{pmatrix} -0.03516 & -0.00719 \\ -0.00305 & -0.03294 \end{pmatrix}$
Drought regime	γ_{01}, γ_{10} (per hour)	0.004013, 0.008036
	$\bar{D}$	1.228
	a_D, b_D^s, b_D^c	-0.098, -0.251, 1.623
Copula	$\zeta(\vartheta) = \tanh(-0.373 + 0.002 s_1 - 0.012 c_1)$	
	$w_C(\vartheta) = (1 + \exp(3.797 + 1.396 s_1 - 3.675 c_1))^{-1}$	
	γ_C	0.0556
Weather-state scales	$\log \varsigma_W(\vartheta) = -0.00004 s_1 + 0.0793 c_1 + d_\sigma(\vartheta)$	
	$\log \varsigma_S(\vartheta) = 0.1020 s_1 - 1.1288 c_1 + 0.0516 s_2 -$	
	$0.1175 c_2$	
Wind map	$\underline{\psi}_W$, $\bar{\psi}_W$	0.00037, 0.9057
	γ_W, c_W	1.1673, 0.9395
	ε_W, δ_W	0.1844, 0.9168
	μ_W^s, μ_W^c	0.1157, 0.6580
Daily wind cycle	$g_{\mu 1}^s$, $g_{\mu 1}^c$, $g_{\mu 2}^s$, $g_{\mu 2}^c$	-0.0676, 0.0748, 0.0358, 0.0287
	e_μ^s, e_μ^c	0.1890, -0.6718
	$g_{\sigma 1}^s$, $g_{\sigma 1}^c$, $g_{\sigma 2}^s$, $g_{\sigma 2}^c$	0.0084, -0.1394, -0.0133, 0.0071
	e_σ^s, e_σ^c	0.0649, -0.2927
Solar map	σ_Γ	0.05977
	κ_0, κ^s, κ^c	1.903, 0.254, -1.208
	κ_2^s, κ_2^c	0.008, -0.053
	ι_0, ι^s, ι^c	0.943, 0.001, -1.125

Table 20: Estimated parameters of the weather model. Harmonics use $s_j = \sin(2\pi j\vartheta/T_Y)$ and $c_j = \cos(2\pi j\vartheta/T_Y)$. The wind-map bounds $\underline{\psi}_W$ and $\bar{\psi}_W$ are set from the observed range of hourly capacity factors; all other parameters are estimated.

E.2 Marginal Density Comparisons

Because the latent state moves gradually toward and away from the drought mean, the hourly marginal distribution is not a mixture of two fixed regime distributions. The model densities in Figure 19 and Figure 20 are therefore estimated from simulated paths of the model over the hours of each month: the stationary density, and the densities over the simulated hours in the normal and in the scarcity regime. In winter, the scarcity regime shifts the wind density markedly toward low capacity factors, whereas the two regime densities nearly coincide in summer. For solar, the regime densities differ less, and the scarcity regime mainly adds mass at low output in winter. The stationary densities follow the data histograms closely in all months.

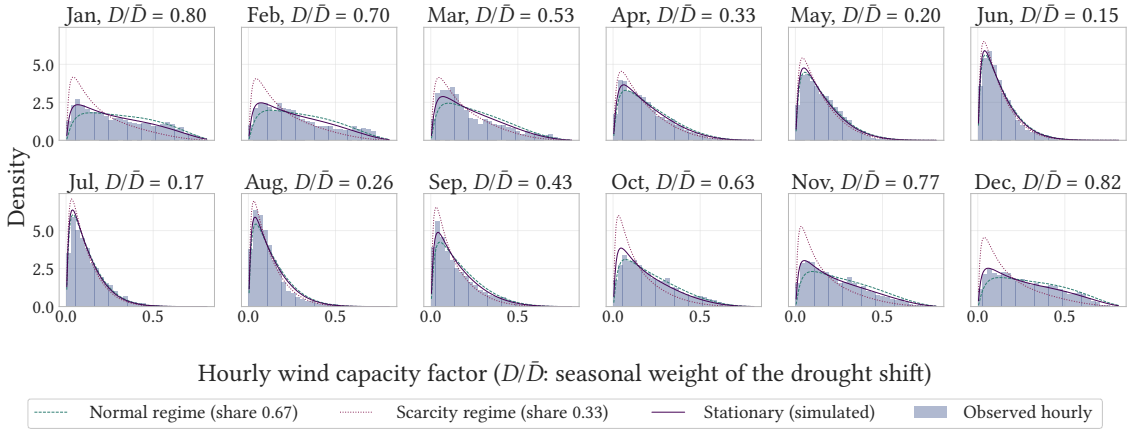


Figure 19: Per-month data histograms and simulated model densities of hourly wind capacity factors: stationary density and densities in the hours of the normal and the scarcity regime. Panel titles give the monthly mean of the seasonal weight $D(\vartheta)/\bar{D}$ of the drought shift.

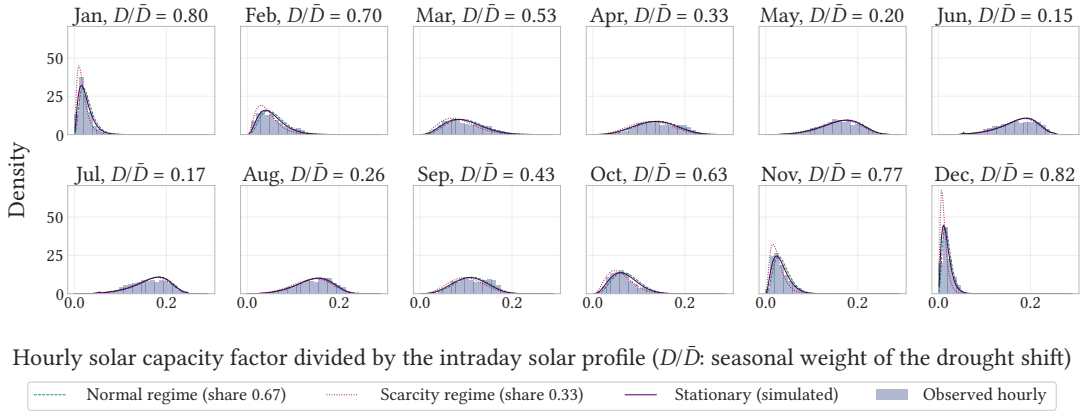


Figure 20: Per-month data histograms and simulated model densities of hourly solar capacity factors divided by the intraday solar profile, in hours in which the profile is at least half of its daily mean: stationary density and densities in the hours of the normal and the scarcity regime. Panel titles give the monthly mean of the seasonal weight $D(\vartheta)/\bar{D}$ of the drought shift.

E.3 Intraday Wind Variation and Drought Days

The daily cycle of the wind map reproduces the timing of calm hours. In Figure 21, the observed share of hours with wind output below 10% of its mean peaks at about 11% around 08:00 UTC outside winter and at about 4.5% around midday in winter. The simulated share peaks at similar times, but its peaks are lower and broader, and in winter about two hours earlier. On a daily basis, the model overstates the share of days on which wind output falls below 10% of its mean (1.3% simulated against 0.6% observed), and it understates the most extreme joint shortfalls of wind and solar: in winter, both daily capacity factors fall below their 5% quantiles on 0.53% of simulated days against 0.76% of observed days. Because the latent state and the drought regime have no memory beyond a few weeks, the model also understates the variation of mean wind production across years. Across simulated 11-year paths, the standard deviation of annual mean wind capacity factors is 6.3% of their mean (median), against 8.4% in the data, and that of winter means 10.6%, against 15.2%; for an equal wind-solar mix, the corresponding values are 4.2% and 9.6%, against 6.0% and 13.4%. These discrepancies may affect the optimal storage capacities in either direction; they are one reason why the comparison with the perfect-foresight optimum on observed weather in Section 6 does not isolate the value of information.

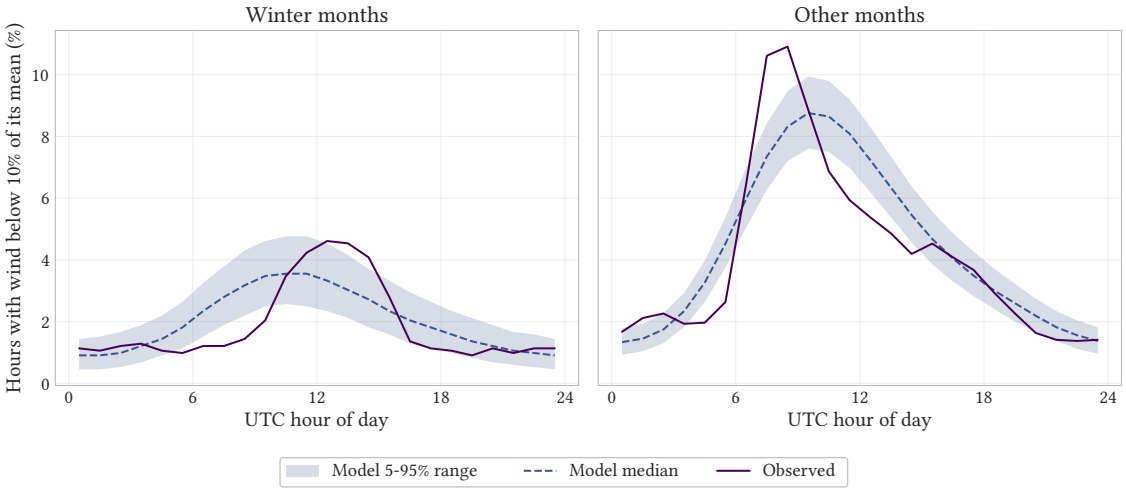


Figure 21: Share of hours with wind output below 10% of its mean, by hour of day (UTC), in winter (November-February) and in the other months: data, and median and 5-95% range across simulated paths.

The quantiles of hourly wind output in Figure 22 show a weaker daily cycle, with a dip in the morning and a rise in the afternoon, which the model reproduces in its median.

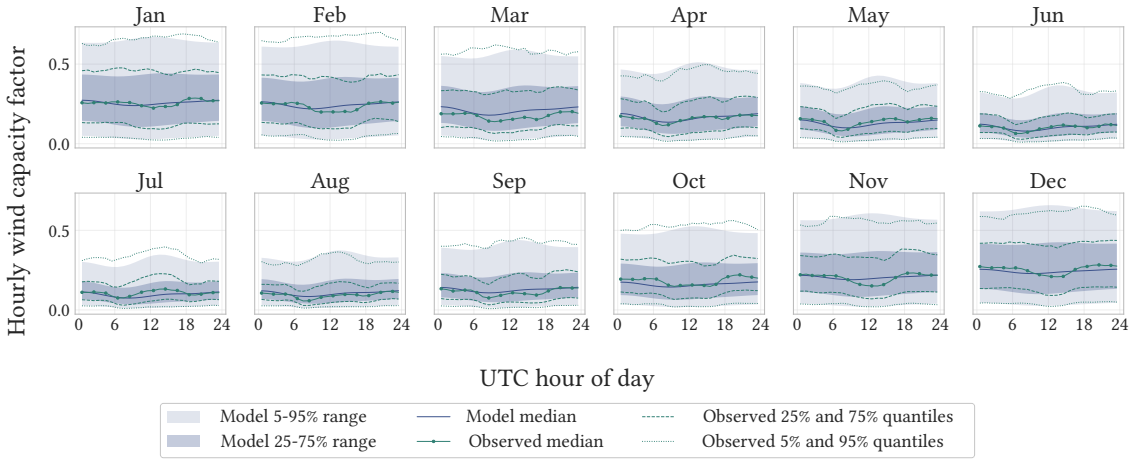


Figure 22: Quantiles of hourly wind capacity factors by hour of day (UTC) for each month: observed median and 25%, 75%, 5%, and 95% quantiles, and simulated median with 25-75% and 5-95% bands.

E.4 Multi-Week Production

Hydrogen storage bridges shortfalls of renewable production that last several weeks. Figure 23 and Figure 3 compare the mean capacity factors of five wind-solar capacity mixes over all overlapping 14- and 28-day windows of the observed record (2015-2025) with those of 8 000 simulated paths of the same length, over the whole year and over the windows that lie entirely in November-February. The simulated distributions follow the data in all panels, including their lower tails.

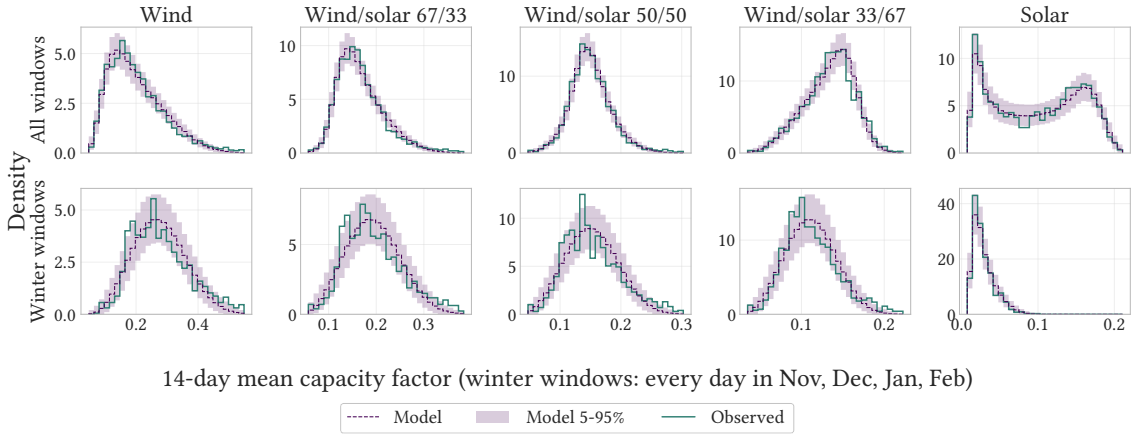


Figure 23: Densities of 14-day mean capacity factors for five wind-solar capacity mixes (columns), over all overlapping windows (top row) and over the windows that lie entirely in November-February (bottom row): observed (2015-2025), model, and the model's 5-95% range across 8 000 simulated paths of the same length.

E.5 Intraday Solar Profile

The intraday solar profile is normalized to a daily mean of one. Let $h = 0, \ldots, N_{\text{slot}} - 1$ index 15-minute slots, with $N_{\text{slot}} = 96$. After removing holidays, for each retained day d, define the empirical ratio of the observed solar capacity factor $\hat{\psi}_{S,d,h}$ in slot h to its daily mean

$$\tilde{\phi}_{S,d,h} = \frac{\hat{\psi}_{S,d,h}}{N_{\text{slot}}^{-1} \sum_j \hat{\psi}_{S,d,j}}. \tag{61}$$

For each slot, estimate a sixth-order Fourier series in the annual phase,

$$\hat{\phi}_S(\vartheta, h) = \beta_{0,h} + \sum_{j=1}^{6} \left(\beta_{j,h}^s \sin\left(\frac{2\pi j \vartheta}{T_Y}\right) + \beta_{j,h}^c \cos\left(\frac{2\pi j \vartheta}{T_Y}\right) \right). \tag{62}$$

The slot-specific fitted values are normalized as

$$\phi_{S,h}(\vartheta) = \frac{\max\left\{0, \hat{\phi}_S(\vartheta, h)\right\}}{N_{\text{slot}}^{-1} \sum_j \max\left\{0, \hat{\phi}_S(\vartheta, j)\right\}}, \tag{63}$$

and, for each hour of day $j = 0, \ldots, 23$, let $\bar{\phi}_{S,j}(\vartheta)$ denote the mean of $\phi_{S,h}(\vartheta)$ over the four slots h of that hour. The profile $\phi_S(\vartheta)$ is the periodic piecewise-linear interpolation, in the time of day $\vartheta \bmod 24$, of the values $\bar{\phi}_{S,j}(\vartheta)$ at the mid-hour points $j + 1/2$. Over the year, it thus inherits the smooth seasonal variation of the Fourier series, and only the time of day is interpolated.

In Figure 24, the simulated quantiles of hourly solar output closely follow the observed ones in every month, both in the timing of production over the day and in its spread across days.

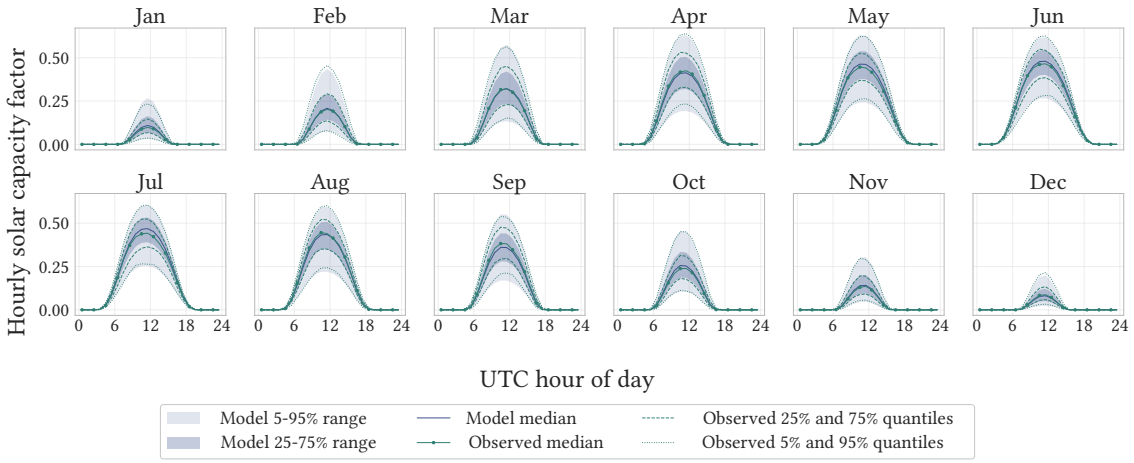


Figure 24: Quantiles of hourly solar capacity factors by hour of day (UTC) for each month: observed median and 25%, 75%, 5%, and 95% quantiles, and simulated median with 25-75% and 5-95% bands.

E.6 Reference Consumption

Let $L_{n,j}$ denote the observed German load, in GW, in quarter-hour $j \in \{0, \ldots, 3\}$ of hour n of 2025 (UTC). With the hourly means $\bar{L}_n = \frac{1}{4} \sum_{j=0}^{3} L_{n,j}, n = 0, \ldots, T_Y - 1$, and the hourly means

p_n of the day-ahead price, the hourly reference consumption is $L_n^{\mathrm{ref}} = \bar{L}_n(\max\{p_n, p_{\mathrm{ref}}\}/p_{\mathrm{ref}})^{|\varepsilon|}$, as in Section 3. Reference consumption is the periodic piecewise-linear interpolation of these values at the mid-hour phases $n + 1/2$:

$$c_{\mathrm{ref}}(\vartheta) = (1-w)L_n^{\mathrm{ref}} + wL_{n+1}^{\mathrm{ref}}, \qquad w = \vartheta - n - 1/2, \quad n + 1/2 \le \vartheta < n + 3/2, \quad (64)$$

with hour indices taken modulo T_Y. Reference consumption thus equals L_n^{ref} of each hour at its midpoint. Since 2025 is not a leap year, its hours match the hours of the model year exactly.

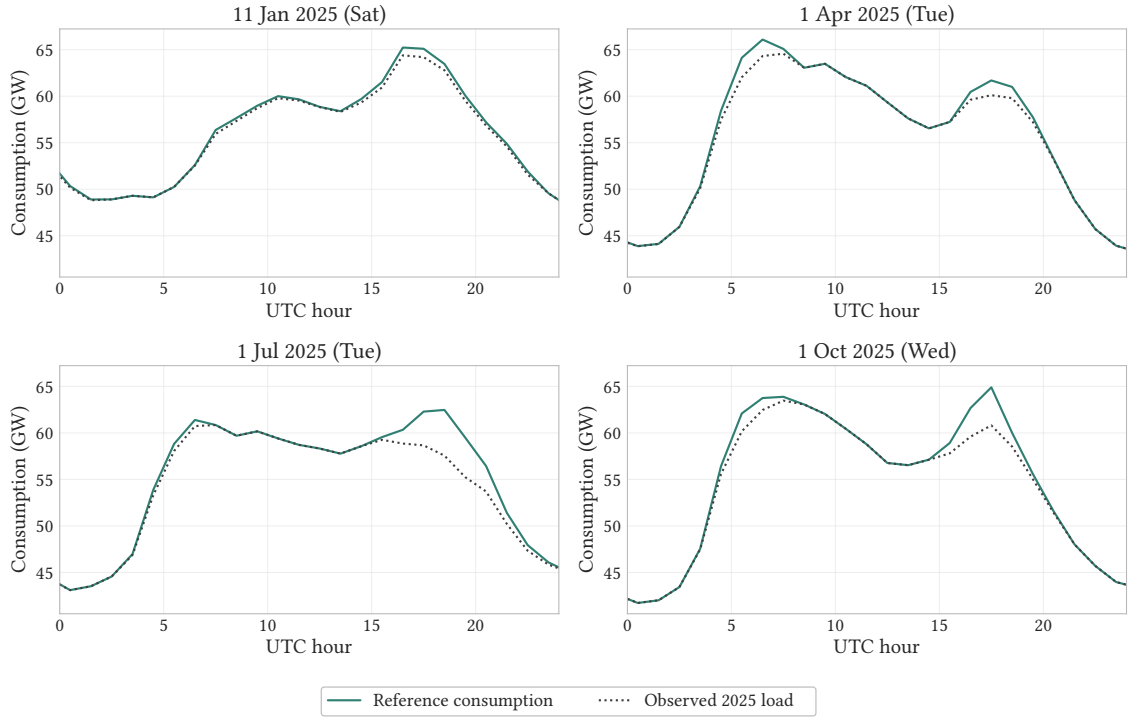


Figure 25: Hourly reference consumption in 2025 for four days, with the observed load (dotted).

Across the selected dates in Figure 25, reference consumption is lowest overnight. On the working days, it rises steeply in the morning and remains elevated through the day, with a dip in the early afternoon and a second peak in the early evening; on the Saturday in January, it rises more gradually and peaks in the late afternoon.

F Overview of the Numerical Solution Method

This appendix summarizes the numerical method used to compute the results. For a fixed capacity vector q, the operating problem is discretized at an hourly time step and solved by a semi-Lagrangian scheme (Chen and Forsyth, 2008; Falcone and Ferretti, 2013), in which storage levels move exactly with the chosen flows and the expected value at the storage levels reached after the hour's decisions is interpolated between storage nodes.

Algorithm summary. For a given capacity vector q, the method proceeds as follows:

1. discretize the phase into the hours of the year, the rotated latent state into a grid for each regime, and the storage levels into a tensor-product grid;
2. sweep backward through the year: at every hour and grid state, choose the storage and fossil actions that maximize the hour's consumption benefit net of fossil cost plus the expected interpolated value of the post-decision state;
3. repeat annual sweeps until the average hourly payoff ρ is stable;
4. propagate the distribution of the discretized state forward under the computed policy to its annual stationary distribution, which yields ρ and its gradient with respect to the capacities;
5. update the capacities with a bounded quasi-Newton step and repeat.

F.1 Hourly Discretization

Hour $n \in \{0, \dots, T_Y - 1\}$ covers the phase interval $[n, n+1)$, with hour 0 starting on January 1 at 00:00 UTC. The storage and fossil actions (α, ν) are chosen at the start of the hour and held for the hour. Stocks then move exactly to $s_i + b_i(\alpha_i)$, with the storage-energy flow b_i of (10) in Section 2. All phase-dependent functions of the hour, including the consumption benefit, are evaluated at the mid-hour phase $\vartheta_n = n + 1/2$. Renewable production of the hour is computed from the weather state at the start of the hour. Write

$$\ell_n(\tilde{x}, r, \alpha, \nu) = U_{\vartheta_n}(y_n(\tilde{x}) + \nu - \alpha_H - \alpha_B) - p_f \nu$$

for the hourly payoff, where $y_n(\tilde{x})$ is renewable production (3) in Section 2 at the rotated latent state $\tilde{x}$ defined below and phase ϑ_n; it does not depend on the regime r.

The weather grid uses the rotated latent coordinates $\tilde{x} = \Omega^T x$, where the columns of the orthogonal matrix Ω are eigenvectors of QQ^T. In these coordinates, the diffusion of the latent Ornstein-Uhlenbeck process (2) of Section 2 is diagonal, its stationary covariance remains the identity, and the drought mean is $-\Omega^T m(\vartheta)$. Each of the two coordinates has 18 uniformly

spaced nodes on $[-5, 5]$, that is, five stationary standard deviations on either side of the mean, and the grid is repeated for each of the two regimes. The one-hour transition matrix $\mathcal{K}_n$ of hour n between the $2 \times 18 \times 18 = 648$ latent states is obtained from the exponential of the joint generator of the regime and the latent state,

$$\mathcal{K}(\vartheta) = \exp\begin{pmatrix} L^0 - \gamma_{01} I, \gamma_{01} I \\ \gamma_{10} I, L^1(\vartheta) - \gamma_{10} I \end{pmatrix}, \tag{65}$$

where $L^r(\vartheta)$ is the generator of the latent process in regime r, discretized with centered differences for drift and diffusion and reflecting ends (Kushner and Dupuis, 2001); in the scarcity regime, its drift points toward the drought mean $-\Omega^T m(\vartheta)$. The matrix exponential is computed exactly at 73 annual phases spaced five days apart, and $\mathcal{K}_n$ interpolates linearly between the two phases around the mid-hour phase ϑ_n.

Both storage grids are denser toward empty and full storage, where the value function bends most. The number of nodes on each storage axis grows with energy capacity, with a target cell size of 340 GWh for hydrogen and 25 GWh for batteries, and lies between 10 and 60 for hydrogen and between 10 and 20 for batteries (60 in the battery-only specification). The baseline two-storage optimum, with about 35 TWh of hydrogen storage and 44 GWh of batteries, has 60 hydrogen and 10 battery nodes, that is, $60 \times 10 \times 648 = 388\,800$ grid states in each of the 8 760 hours; the battery-only baseline optimum has 30 battery nodes.

F.2 Bellman Recursion and the Average Payoff

Write w_n for the relative value at the start of hour n on the grid, periodic in the year, $w_{T_Y} = w_0$. With the regime observed, the hourly Bellman equation reads

$$w_n(\tilde{x}, r, s) = \max_{(\alpha,\nu)\in A_n(\tilde{x},r,s)} \left\{ \ell_n(\tilde{x}, r, \alpha, \nu) + \sum_{(\tilde{x}',r')} \mathcal{K}_{n,(\tilde{x},r),(\tilde{x}',r')} w_{n+1}(\tilde{x}', r', s + b(\alpha)) \right\} \tag{66}$$

where $A_n(\tilde{x}, r, s)$ collects the actions that respect the storage bounds and power limits, fossil capacity, and nonnegative consumption, as in (37) of Appendix B. Post-decision stocks $s + b(\alpha)$ generally lie between grid nodes. Along each storage axis, the continuation value is represented by a piecewise-cubic interpolant of its values at the grid nodes, which preserves the monotonicity of the value in the stock (Fritsch and Butland, 1984). The maximization over actions is exact for a piecewise-linear approximation of this interpolated continuation value and is described in Appendix G.

The scheme solves (66) by relative value iteration (Puterman, 1994). Starting from $w = 0$, each annual sweep applies the hourly maximization backward from hour $T_Y - 1$ to hour 0

and returns the value $\tilde{w}_0$ at the start of the year. The statewise annual increments $\Delta = (\tilde{w}_0 - w_0)/T_Y$ become equal across states as the iteration converges, their mean is the current estimate of ρ, and the value is renormalized to zero mean before the next sweep. The iteration stops when the spread $\max \Delta - \min \Delta$ falls below $10^{-7}|\rho|$, or after at most 60 sweeps. In every evaluation underlying the results, the iteration converged before reaching this cap. Started from the value of the previous capacity candidate, the evaluation of the baseline optimum converges in three sweeps. The annual operating value is $T_Y \rho$, as in (12) of Section 2.

The perfect-foresight problems of Appendix M use the same scheme with a single latent state: the Bellman recursion runs over all hours of the period and evaluates production (70) in Appendix M at the hour midpoints, where it equals the hourly production of the path.

F.3 Stationary Distribution and Capacity Gradient

At the computed policy, the recursion is linear: $w_n = \ell^q_n - \rho \mathbf{1} + P_n w_{n+1}$, where ℓ^q_n is the vector of optimal hourly payoffs and P_n maps each grid state to the interpolation nodes around its post-decision stocks and then through $\mathcal{K}_n$. The distribution of the discretized state at the start of the year is the stationary distribution μ_0 of the annual map, $\mu_0 = \mu_0 P_0 P_1 \cdots P_{T_Y - 1}$, and the hourly distributions follow from $\mu_{n+1} = \mu_n P_n$. Then

$$\rho = \frac{1}{T_Y} \sum_{n=0}^{T_Y - 1} \mu_n \cdot \ell^q_n, \qquad \frac{\partial \rho}{\partial q_j} = \frac{1}{T_Y} \sum_{n=0}^{T_Y - 1} \mu_n \cdot \frac{\partial h_n}{\partial q_j}, \tag{67}$$

where h_n is the hourly maximum in (66), differentiated with respect to capacity q_j at fixed continuation values (Cao, 2007). This is the discrete counterpart of the envelope formula in Appendix C. The distribution is propagated forward through the year (Young, 2010) until its change over a year falls below 10^{-8} in total variation, and the gradient is accumulated over the final year. Because $\mathcal{K}_n$ and the cubic interpolation weights contain small negative entries, μ_n is a signed measure of unit mass with negligible negative part.

F.4 Simulation of the Optimal Operation

All consumption, price, revenue, and welfare measures in the results, including NPV and the levelized costs, are long-run averages (16) in Section 2 estimated by simulating the computed policy. Each of 4 096 paths starts from a state drawn from the stationary distribution μ_0, spread continuously between neighboring grid nodes, and then runs at hourly resolution with continuous storage levels. Weather is simulated without the grid: the latent state follows its exact hourly autoregressive transition toward the target of the regime at the start of the hour, and

the regime its exact hourly transition probabilities. After two burn-in years, 20 years per path are recorded. The simulated policy respects the storage bounds, power limits, fossil capacity, and nonnegative consumption in every hour, so its simulated payoff is attained by a feasible policy. The simulated weather transition holds the regime fixed within the hour, whereas the transition used in value iteration lets the regime switch within the hour. With hourly switching probabilities of about 0.4% from the normal and 0.8% from the scarcity regime, this difference is likely small. Up to Monte Carlo error and this difference, the simulated payoff is therefore a lower bound on the optimal long-run payoff of the operating problem with hourly decisions. Appendix H compares it with the value-iteration payoffs; their difference combines the effects of the grid, the interpolation, and the weather transition. The capacity optimization itself uses the value-iteration payoff ρ rather than these simulations.

F.5 Residual of the Continuous-Time HJB

A direct check of accuracy measures how far the computed solution is from satisfying the continuous-time HJB (39) in Appendix B. For the baseline two-storage optimum, the hourly relative values w_n are interpolated across hours, storage nodes, and weather nodes into a function $\widetilde{W}$ that is continuously differentiable in the phase and the stocks and twice continuously differentiable in the latent state, the smoothness required by the verification argument in Appendix B. The residual of the HJB at state ξ is

$$\mathcal{E}(\xi) = \sup_{(\alpha,\nu)\in A(\xi)} \left\{ \ell^q(\xi,\alpha,\nu) + \mathcal{L}^{\alpha,\nu}\widetilde{W}(\xi) \right\} - \rho,$$

where ρ is the value-iteration payoff. It is the amount by which the best instantaneous return at ξ, with stored energy valued at the stock derivatives of $\widetilde{W}$, exceeds the average hourly payoff, and it vanishes at every state for an exact solution.

The residual is evaluated along one recorded year of the 4 096 stationary simulated paths, at a time drawn uniformly within each path-hour, that is, at about 35.9 million states. At that time, the stocks have moved at the flows of the hour, the latent state has followed its exact transition under the regime of the hour, and reference consumption and the intraday solar profile are interpolated linearly between hour midpoints. Averaged over these states, the residual is 0.42 B€/y, or 0.90 €/MWh of mean reference consumption, with a standard error of 0.016 €/MWh computed from the annual means of the paths. In present value, it amounts to 8.5 B€, about 0.6% of the absolute NPV, and per MWh of reference consumption to about 0.6% of the welfare-adjusted LCOE from value iteration on the baseline grid, 142.43 €/MWh (Table 21). Its positive sign agrees with the direction of the refined-grid evaluation in Appendix H, which raises the

operating value by 0.10 B€/y. Across states, the median residual is 0.15 €/MWh and the mean absolute residual 3.85 €/MWh; 90% of the residuals lie between −6.47 and 10.10 €/MWh. The largest residuals occur at low hydrogen fill, where the value function bends most: states with less than 20% hydrogen fill, 1.2% of the sample, have a mean residual of 14.9 €/MWh, against 0.43 €/MWh for states above 80% fill, which make up 63% of the sample.

F.6 Capacity Optimization

In the restriction without fossil backup, the free components of the capacity vector q are

$$\left(K_W, K_S, \bar{S}_H, \bar{S}_B, K_{H-}, K_{H+}\right),$$

with the notation of Section 2. Battery power capacities follow from the two-hour duration assumption in Section 4, and fossil capacity is zero. I impose optimization bounds: $K_W, K_S \in [100, 500]$, $\bar{S}_H \in [1, 100\,000]$, $\bar{S}_B \in [5, 5\,000]$, $K_{H-} \in [1, 100]$, and $K_{H+} \in [1, 100]$, with power in GW and energy in GWh. The battery-only specification uses $K_W, K_S \in [100, 1\,000]$ with the same battery bound, and the fossil-backup specification adds $F \in [0.1, 80]$ and lowers the lower battery bound to 0.5 GWh. None of these bounds binds at the optima, except the lower bounds on hydrogen capacities at low carbon prices in Section 7. In these cases, the restricted mix without hydrogen storage is optimized separately, and the mix with the higher NPV is selected.

Each capacity is mapped to the unit interval: linearly for generation and power capacities and logarithmically for the two storage energy capacities, so that a fixed step corresponds to a fixed relative change in energy capacity. The objective is the net present value NPV^q in (13) of Section 2, maximized with the bounded quasi-Newton method L-BFGS-B (Byrd et al., 1995), using the gradient (67) and the gradient of the linear capacity cost $C(q)$.

The search proceeds in stages of increasing resolution. The first stage searches the whole capacity box on a coarse grid with 10 nodes per storage axis and 15×15 latent nodes per regime. The later stages set the storage grid at the previous optimum by the rule for the number of storage nodes given above, use the 18×18 latent grid, and search a reduced box around that optimum, in which each capacity can move by 30% of its value or, for small capacities, by a minimum amount; the final stage also refines the hourly maximization as described in Appendix G. Further stages follow if the optimum lies on the boundary of its reduced box. A stage stops when the relative improvement of the NPV falls below a tolerance, 10^{-7} in the final stage, and each evaluation starts from the value function and the stationary distribution of the previous candidate. For the baseline optimum, the NPV computed on the three stage grids differs by less than 0.5%. The optimal capacities also change little between the coarse grid of the first

stage and the final grid: wind 308 and 312 GW, solar 224 and 216 GW, hydrogen storage 34.9 and 35.0 TWh, battery storage 43 and 44 GWh, electrolyzer 48 and 47 GW, and hydrogen CCGT 46 and 45 GW.

One capacity optimization runs on a single NVIDIA A100 GPU. At the final-stage grid of the baseline optimum, one evaluation of ρ and its gradient takes about six minutes, and the complete optimization about 80 minutes.

F.7 Cost Elasticities

Differentiating the investment conditions (14) in Section 2 with respect to the unit costs gives $T_Y \nabla^2 \rho^{q^*} dq^* = \delta dC$, where $\nabla^2 \rho^{q^*}$ is the matrix of second derivatives of ρ^q with respect to the independently chosen capacities at the optimum. The cost elasticities (23) in Section 5.1 are therefore

$$e_{ij} = \frac{\delta}{T_Y} \left[\left(\nabla^2 \rho^{q^*} \right)^{-1} \right]_{ij} \frac{C_j}{q_i^*}. \tag{68}$$

If this matrix is negative definite, every own-cost elasticity e_{ii} is negative, and e_{ij} and e_{ji} have the same sign.

The matrix $\nabla^2 \rho^{q^*}$ is computed by central differences of the capacity gradient (67): each free capacity is moved to $q_j^*(1 \pm 0.01)$, the other capacities fixed, and column j of the matrix is the difference of the two gradients divided by $0.02 q_j^*$. The evaluations, two for each free capacity and one at the optimum, use the final-stage grid of the optimum, each starting from the value function and the stationary distribution of the previous one, and each cross derivative is the average of its two estimates, $\partial_{q_i} \partial_{q_j} \rho$ and $\partial_{q_j} \partial_{q_i} \rho$. At the three battery-only optima of Section 4 and the three two-storage optima of Section 5, these estimates differ by at most 0.12% of the geometric mean of the corresponding own second derivatives. The resulting matrix is negative definite at each optimum, so the second-order condition of the capacity problem holds.

G Optimal Policy with Two Storage Technologies

This appendix describes the hourly maximization in (66) of Appendix F. At every hour and grid state, the storage actions $\alpha = (\alpha_H, \alpha_B)$ and fossil dispatch ν maximize the payoff of the hour plus the continuation value $\bar{w}$ of the post-decision stocks $s + b(\alpha)$, where $\bar{w}(s') = \sum_{(\tilde{x}', r')} \mathcal{K}_{n,(\tilde{x},r),(\tilde{x}',r')} w_{n+1}(\tilde{x}', r', s')$ is the expected continuation value in (66) for the belief-based policy of Appendix N, it is the belief-weighted average of the continuation values of the two regimes.

G.1 Piecewise-Linear Continuation Value

Within one hour, stock i can reach the interval between $s_i - \omega_{i-} K_{i-}$ and $s_i + \omega_{i+} K_{i+}$, intersected with $\left[0, \bar{S}_i\right]$. This reach is split into pieces at the current stock, at the grid nodes, and at $M - 1$ equally spaced points between the current stock and each end of the reach. On each piece, the post-decision stock is linear in the action, and the monotone cubic interpolant of the continuation value is replaced by the straight line through its values at the two ends of the piece. Each pair of pieces, one per storage technology, is cut along its diagonal into two triangles (Kuhn, 1960) on which $\bar{w}$ is linear in the actions; with nonnegative consumption, each triangle defines a polygon of feasible actions. Replacing the cubic by this line on a piece of length l changes the continuation value by at most $l^2 \max|\bar{w}''|/8$. The final optimization stage and the simulations of the optimal operation use $M = 4$ for both storage technologies; the earlier stages use $M = 1$.

G.2 Exact Maximization

On each polygon, the objective is the concave consumption benefit plus a linear function of the actions. It is linear along the direction $(1, -1)$, which leaves consumption unchanged, so its maximum lies on an edge. On an edge, the marginal value of consumption equals the gain in continuation value per unit of consumption forgone, a condition solved in closed form with the explicit inverse of the marginal value from Appendix A. The objective is not concave across polygons. The objective on each polygon is therefore bounded from above by replacing the consumption benefit with a tangent, and polygons are solved in decreasing order of these bounds, starting from the polygon that was optimal at the same state in the previous sweep, until no remaining bound exceeds the best value found. The resulting action is the exact maximizer of the objective with the piecewise-linear approximation of the interpolated continuation value up to rounding error; at the baseline optimum, about 1.2 polygons per state are solved on average.

G.3 Fossil Dispatch

With fossil capacity, the consumption benefit of the hour is replaced by its maximum over fossil dispatch, $U^F(c_0, F_0) = \max_{0 \le \nu \le F_0} \left\{ U_{\vartheta_n}(c_0 + \nu) - p_f \nu \right\}$ with $c_0 = y - \alpha_H - \alpha_B$, which is concave and attained by the dispatch rule of Section 2 with available fossil capacity F_0. With independent fossil capacity, $F_0 = F$. With the hydrogen CCGTs of Section 7, gas and hydrogen share the turbine capacity, and $F_0 = F + K_{H-} + \min\{\alpha_H, 0\}$ shrinks one-for-one with hydrogen discharge. On each edge, the maximization then splits into at most three closed-form dispatch cases, and the bounds use a supergradient of U^F in place of the tangent of U.

H Robustness to the State Grid Size

This appendix examines the robustness to the grid of weather and storage states of the fully renewable two-storage optimum with baseline technology costs, a 5% discount rate, and $|\varepsilon| = 0.05$. In each of the 8 760 hours of the year, the optimization uses 18×18 latent nodes for each of the two regimes, 60 hydrogen-storage nodes, and 10 battery-storage nodes, that is, 388 800 grid states per hour. The robustness exercise keeps the hourly time step and doubles the number of nodes along each latent and storage dimension, giving $36 \times 36 \times 2 \times 120 \times 20 = 6\,220\,800$ grid states per hour.

The finer grid is used to re-evaluate the optimum by relative value iteration (Table 21).

Grid	Grid states per hour	Operating value [B€/y]	NPV [T€]	$\mathrm{LCOE}_{\mathrm{adj}}$ [€/MWh]
Baseline grid, simulated	388 800	−2.097	−1.333772	142.19
Baseline grid, value iteration	388 800	−2.211	−1.336065	142.43
Refined grid, value iteration	6 220 800	−2.114	−1.334124	142.22

Table 21: Evaluation of the optimum on the baseline and refined grids. The simulated row is computed by simulating the policy, as are the values in other tables throughout the paper. The value-iteration rows come directly from relative value iteration on the respective grid.

Relative to the value-iteration evaluation on the baseline grid, the refined-grid evaluation raises NPV by 1.94 B€, about 0.15% of its absolute level, and lowers $\mathrm{LCOE}_{\mathrm{adj}}$ by 0.21 €/MWh. The annual operating value rises by 0.10 B€/y, which is small relative to the annualized capacity cost of about 65 B€/y. This discretization effect is small relative to the LCOE differences between the cases in Section 5.

I Results with Low Battery Costs

This appendix shows the capacities and costs, the operation, and the consumption and price distributions of the optimal mix with battery and hydrogen storage under the low-battery-cost assumption at a 5% discount rate, the counterparts of the baseline figures in Section 5.

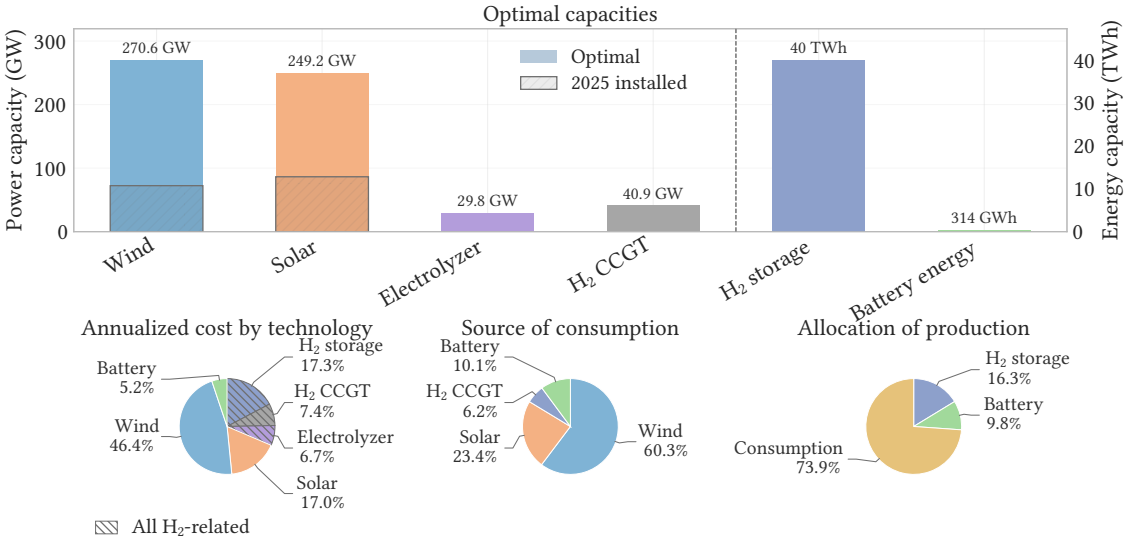


Figure 26: Optimal and 2025 installed capacities, annualized costs by technology, sources of consumed electricity, and allocation of produced electricity for the low-battery-cost case at a 5% discount rate. The 2025 installed wind capacity comprises onshore and offshore wind.

Figure 27, Figure 28, and Figure 29 illustrate the operation of the optimal mix under the low-battery-cost assumption. Along the simulated path, the hydrogen level falls from about 32 TWh in early January to about 7 TWh at the beginning of March and is rebuilt over the summer. In the December week, hydrogen is discharged at the full 41 GW of CCGT capacity on most days, and the battery runs empty.

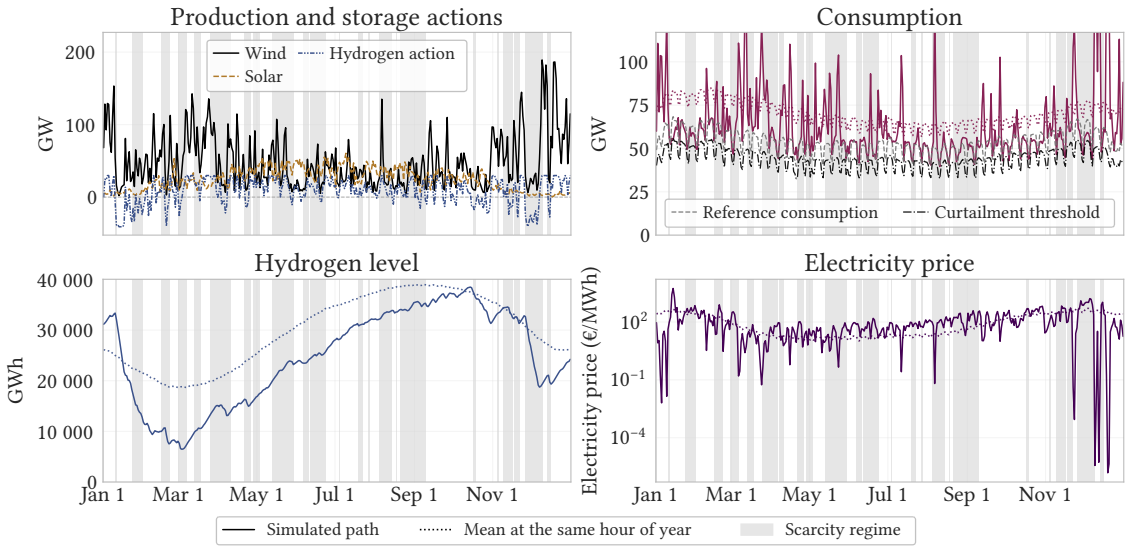


Figure 27: Simulated year beginning in January for the low-battery-cost case at a 5% discount rate: daily means along one simulated path, with the hydrogen storage level. The mean at the same hour of year is taken over all 4 096 simulated paths of 20 years each.

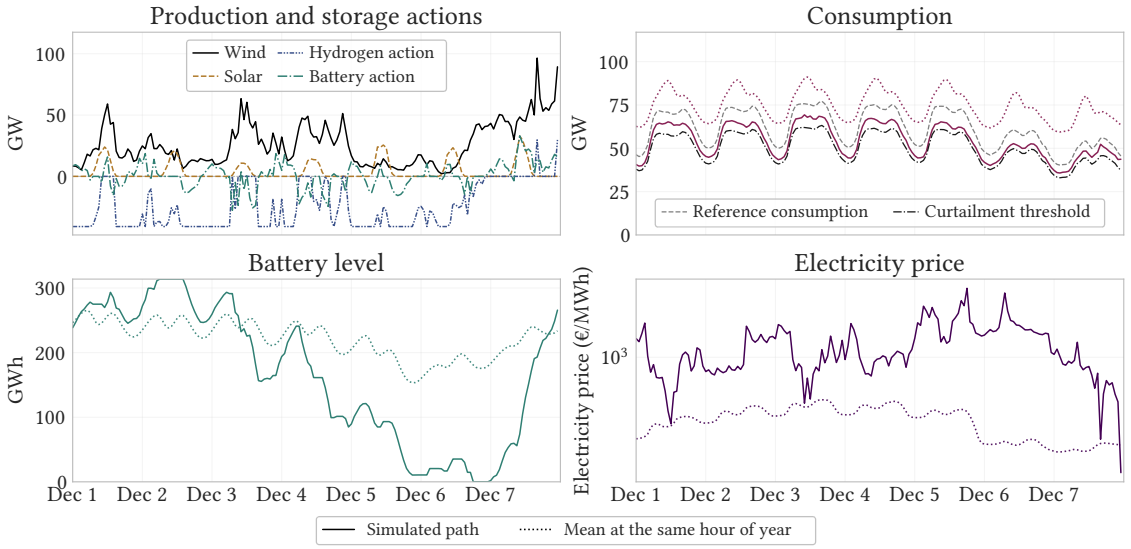


Figure 28: Simulated week in December for the low-battery-cost case at a 5% discount rate: hourly values along one simulated path, with the battery storage level. The mean at the same hour of year is taken over all 4 096 simulated paths of 20 years each.

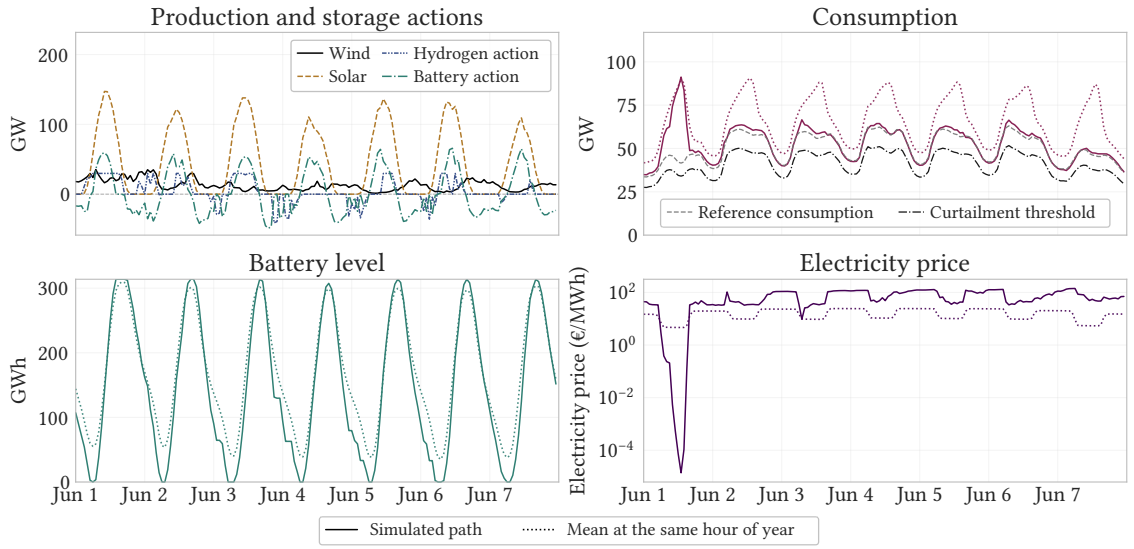


Figure 29: Simulated week in June for the low-battery-cost case at a 5% discount rate: hourly values along one simulated path, with the battery storage level. The mean at the same hour of year is taken over all 4 096 simulated paths of 20 years each.

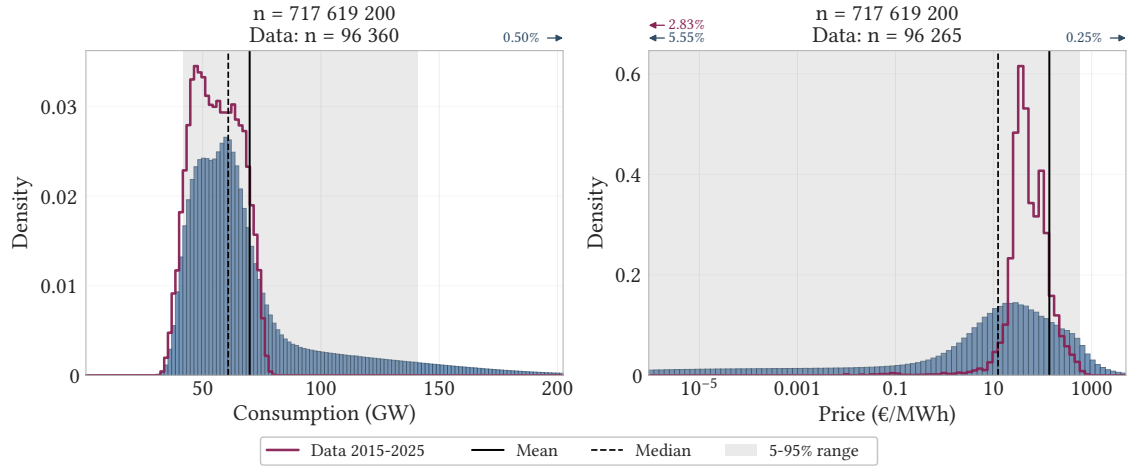


Figure 30: Simulated consumption and shadow-price histograms for the low-battery-cost case at a 5% discount rate, with the observed hourly load and day-ahead prices for 2015-2025 (Data). The mean, median, and 5-95% range refer to the simulated values; the arrows give the shares of simulated (lower arrows) and observed (upper arrows) values beyond the plotted range, and n counts hours.

J Decentralization by Supporting Electricity Prices

This appendix states conditions under which the planner's allocation in Section 2 can be decentralized by supporting electricity prices. Consumers and firms are price takers, all agents have the same information as the planner, and they maximize risk-neutral expected payoffs. It also compares simulated revenues with lifecycle costs at the optimum and reports the average sale prices at the optimum with hydrogen storage only.

J.1 Convex Formulation and Assumptions

Use the notation and technologies of Section 2, with the linear lifecycle capacity cost $C(q) = \sum_j C_j q_j$ of (47) in Appendix D. Multiplying the planner's objective by δ/T_Y leaves its maximizers unchanged, so investment and operation jointly maximize long-run hourly consumption benefits net of fossil costs and the hourly capital charge $\delta C(q)/T_Y$. For the convexity argument, let the charging and discharging components $a_{it+} \geq 0$ and $a_{it-} \leq 0$ be chosen independently within their power limits. All constraints are then linear in capacities and operating paths, convex combinations along the same weather histories remain adapted to the available information, and the planner maximizes a concave objective over a convex set. This relaxation has the same optimal value: replacing simultaneous charging and discharging by the net flow leaves the energy flow into storage unchanged and, because $\omega_{i+} \leq \omega_{i-}$, does not reduce the electricity available for consumption. An optimum can therefore be chosen with $a_{it+}a_{it-} = 0$, so that $a_{it} = a_{it+} + a_{it-}$ is the signed action of Section 2.

Assume further that the expected time averages of benefits, costs, and electricity payments exist and are finite for the allocations and deviations considered, and that the planner problem satisfies strong duality with an attained dual optimum, an adapted electricity-price process P^*: after energy balance is relaxed at P^*, the maximal value of the Lagrangian equals the planner's value. Convexity alone does not establish this supporting-multiplier condition in an infinite-horizon problem.

J.2 Supporting Prices

Proposition 1 (Decentralization). *Under these assumptions, the planner's capacities q^* and operation, together with the electricity prices $P_t^* = U'_{\theta_t}(c_t^*) = p_t^{q^*}$, where $c_t^* = c_t^{q^*}$, form a competitive equilibrium in which price-taking consumers and firms maximize risk-neutral expected payoffs. Conversely, every market-clearing competitive allocation at these prices is planner-optimal.*

Proof sketch. Attach the multiplier P_t^* to energy balance in supply-minus-demand form, retaining all physical and information constraints within each technology. The hourly Lagrangian separates as

$$\begin{aligned}\mathcal{G}_t &= U_{\theta_t}(c_t) - p_f f_t - \delta C(q)/T_Y + P_t^* \left(Y_t + f_t - c_t - \sum_i a_{it} \right) \\ &= \left[U_{\theta_t}(c_t) - P_t^* c_t \right] + \sum_{j \in \{W,S\}} P_t^* y_{jt} + \left(P_t^* - p_f \right) f_t \\ &\quad - \sum_i P_t^* a_{it} - \delta C(q)/T_Y,\end{aligned} \tag{69}$$

where y_{jt} is the output of renewable technology j, so that $Y_t = y_{Wt} + y_{St}$. The terms of $\mathcal{G}_t$ are consumer surplus, renewable revenue, the fossil operating margin, storage revenue net of charging expenditure, and the capital charge; shared equipment, such as the hydrogen CCGTs, keeps its joint operating constraints and is optimized as one production activity. At the planner's allocation, the balance term vanishes, so dual optimality implies that this allocation, including its capacities, maximizes the long-run average of $\mathcal{G}_t$ with energy balance relaxed. Each consumption, operating, or investment choice then enters only its own payoff and feasible set. If a consumer or firm could raise its payoff at the fixed process P^*, by changing its operation, its capacity, or both, substituting that choice alone would raise the Lagrangian. The planner's choices are therefore individually optimal, and their energy balance clears the market. Conversely, at a market-clearing competitive allocation, individual optimality makes the sum of payoff differences relative to any feasible alternative nonpositive. Electricity payments cancel because energy balance holds in both allocations, so no feasible alternative raises planner welfare. □

Price taking means that a deviation leaves the whole stochastic process P^* fixed. The process can vary with the aggregate state and the observed weather history without revealing future weather, and a storage operator maximizes the long-run average of $-P_t^* a_{it}$ subject to its stock and power constraints. On the curtailment branch, $c_t^* < k(\theta_t)$, the supporting price exceeds p_{cap}, so an externally imposed cap would generally remove the supporting incentives. At a differentiable interior optimum, the envelope argument of Appendix C values an extra unit of equipment identically for the firm and the planner, through the rents on the constraints it relaxes, so the investment condition (14) in Section 2 holds in both problems.

J.3 Revenues and Lifecycle Costs at the Optimum

At the supporting prices, the revenue of each technology is its operating margin, defined in Appendix D. Table 22 compares annual revenues with annualized lifecycle costs for the optima with battery and hydrogen storage in Section 5. Revenues are within about 3% of annualized lifecycle costs for every technology, as the investment conditions (14) in Section 2, with the marginal operating benefits of Appendix C, require at an interior optimum. The remaining gaps, with revenues slightly below costs in every entry, are small and arise because revenues are estimated by simulating the computed policy (Appendix F), whereas the capacity optimization uses the value-iteration payoff, which differs slightly from the simulated payoff (Table 21).

Case	Wind revenue [B€/y]	Solar revenue [B€/y]	H_2 storage revenue [B€/y]	Battery storage revenue [B€/y]
Baseline, 5%	32.51 (33.01)	8.90 (9.11)	20.62 (20.90)	1.53 (1.56)
Baseline, 2%	22.16 (22.55)	6.07 (6.22)	14.64 (14.84)	1.11 (1.13)
Low-battery-cost, 5%	28.37 (28.62)	10.27 (10.50)	19.19 (19.36)	3.08 (3.18)

Table 22: Revenue summary; annualized lifecycle costs in parentheses. Hydrogen storage revenue and costs cover cavern storage, electrolyzers, and hydrogen CCGTs together.

J.4 Average Sale Prices with Hydrogen Storage Only

Table 23 shows the quantity-weighted average sale prices (49) in Appendix D for the optimum with hydrogen storage only (Section 5.5).

Case	Wind [€/MWh]	Solar [€/MWh]	H_2 storage [€/MWh]
Baseline, 5%	59.34	47.96	347.95
Baseline, 2%	40.05	30.47	244.14

Table 23: Quantity-weighted average sale prices with hydrogen storage only: wind and solar per MWh produced; storage as net revenue, discharge revenue less charging expenditure, per MWh discharged. Hydrogen storage covers cavern storage, electrolyzers, and hydrogen CCGTs together.

K NPV Sensitivity and Cost Elasticities

This appendix evaluates how NPV changes when each capacity is varied around the optimum with baseline technology costs and a 5% discount rate, with $|\varepsilon| = 0.05$, two storage technologies, and no fossil backup. The reference vector is the baseline row at a 5% discount rate in Table 7.

Each capacity is varied separately from 10% below to 10% above its optimized value in one-percentage-point increments, while all other capacities remain fixed. Figure 31 shows the change relative to the NPV at the computed optimum.

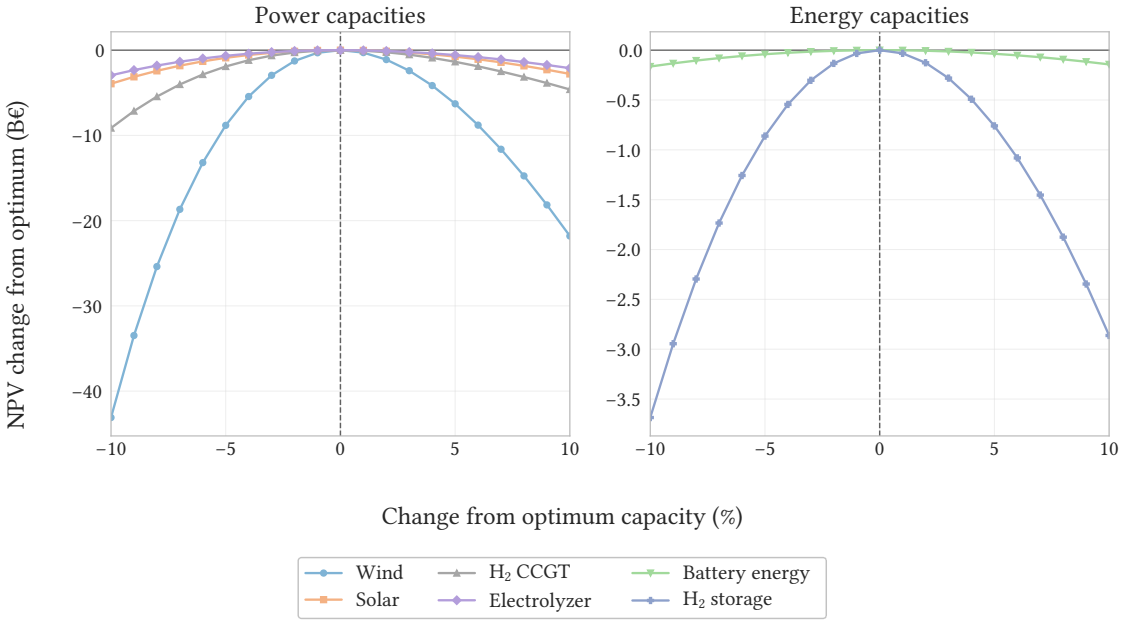


Figure 31: NPV changes from varying one capacity around the baseline optimum at a 5% discount rate, computed by relative value iteration. Other capacities are fixed.

NPV declines at both endpoints of the range for every capacity. Wind deviations are the most costly, lowering NPV by about 43 B€ at 10% below and 22 B€ at 10% above the optimum, followed by the hydrogen CCGT, with 9 and 5 B€. Solar, hydrogen storage, and electrolyzer capacities follow with changes between 2 and 4 B€, and battery energy capacity is the flattest direction, with changes of about 0.15 B€. Reductions in capacity have a larger effect than equal-sized increases for every capacity. At the center of each variation, the independently recomputed NPV differs from the value obtained in the capacity optimization by less than 0.001 B€, and no point of these variations exceeds the NPV at the center by more than 0.001 B€.

The flat battery direction is reflected in the cost elasticities of Table 10: with little curvature of NPV in battery capacity, battery capacity responds most strongly to its own cost, with an

elasticity of −1.95. Table 24 reports the cost elasticities at the battery-only optima. In all three cases, wind, solar, and battery capacities substitute for one another.

Case	Capacity	Unit cost of		
		Wind	Solar	Battery
Baseline, 5%	Wind	−0.733	0.353	0.196
	Solar	0.808	−1.269	0.280
	Battery	0.438	0.275	−1.142
Baseline, 2%	Wind	−0.758	0.374	0.214
	Solar	0.813	−1.258	0.256
	Battery	0.451	0.248	−1.099
Low-battery-cost, 5%	Wind	−0.599	0.286	0.172
	Solar	0.916	−1.560	0.550
	Battery	0.296	0.296	−0.867

Table 24: Cost elasticities e_{ij} of the optimal capacities with battery storage only: the percentage change in the optimal capacity of the row when the unit lifecycle cost of the column rises by 1%.

L Robustness to Curtailment-Cost, Renewal-Cost, and Reference-Consumption Assumptions

This appendix re-optimizes the fully renewable mix with battery and hydrogen storage at baseline technology costs, a 5% discount rate, and $|\varepsilon| = 0.05$, changing one assumption at a time. The first two cases set the blackout value v_{blackout}, the marginal value of electricity at zero consumption, to 25 000 and 100 000 €/MWh instead of 50 000 €/MWh. The next two set the price cap p_{cap} to 2 500 and 10 000 €/MWh instead of 5 000 €/MWh. Since the elasticity is held fixed, the price cap also moves the curtailment threshold $k(\vartheta) = \kappa c_{\text{ref}}(\vartheta)$ of Appendix A: κ equals 0.848 and 0.791 (Table 25), so the share of price-sensitive demand is 15.2% and 20.9%, compared with 18.1% in the baseline. The fifth case sets the refurbishment fraction to $\zeta_j = 1$ for every technology in (47) of Appendix D, so that the full initial CAPEX is paid again at the end of each lifetime. The last two cases change the construction of reference consumption in Section 3. The first uses the observed 2025 load without the price correction. The second applies the correction in both directions, so that reference consumption also lies below the observed load in hours priced below p_{ref}, with prices below 1 €/MWh set to 1 €/MWh. Each case is optimized by the procedure of Appendix F and simulated on weather paths generated from the same random draws as the baseline.

Case	κ	Wind [GW]	Solar [GW]	H_2 storage [GWh]	Battery [GWh]	Electrolyzer [GW]	H_2 CCGT [GW]
Baseline	0.819	312	216	35 010	44	47	45
$v_{\text{blackout}} = 25\,000$	0.819	311	215	33 226	44	47	45
$v_{\text{blackout}} = 100\,000$	0.819	313	217	37 002	45	47	46
$p_{\text{cap}} = 2\,500$	0.848	312	216	34 685	44	47	45
$p_{\text{cap}} = 10\,000$	0.791	313	216	35 756	45	47	46
$\zeta_j = 1$	0.819	305	223	36 204	38	48	46
c_{ref} = observed load	0.819	309	215	34 669	42	46	45
Two-sided price correction	0.819	310	207	34 777	44	47	46

Table 25: Optimized capacities for the curtailment-cost, renewal-cost, and reference-consumption cases (baseline technology costs, 5% discount rate, $|\varepsilon| = 0.05$). Blackout values and price caps in €/MWh; κ is the curtailment threshold relative to reference consumption.

The curtailment-cost parameters mainly affect hydrogen storage capacity (Table 25). Halving the blackout value lowers it from 35.0 to 33.2 TWh, and doubling the blackout value raises it to 37.0 TWh. A higher cost of curtailment raises the value of a reserve for long low-production periods, which cavern storage provides at the lowest cost. The welfare-adjusted LCOE lies

between 141.12 and 143.40 €/MWh in the four curtailment-cost cases, compared with 142.19 €/MWh in the baseline (Table 26).

Case	$\bar{c}_{\text{ref}}$ [GW]	$\bar{c}^q$ [GW]	SD (c^q) [GW]	$\bar{c}^q_{\text{CE}}$ [GW]	$\bar{p}^q$ [€/MWh]	$\delta\text{Cost}^q_{\text{sys}}$ [B€/y]	$\delta\text{Cost}^q_{\text{tot}}$ [B€/y]	LCOE [€/MWh]	LCOE_{CE} [€/MWh]	LCOE_{adj} [€/MWh]
Baseline	53.54	71.23	34.71	51.74	144.81	64.59	66.69	103.52	142.52	142.19
$v_{\text{blackout}} = 25\,000$	53.54	71.00	34.54	51.66	143.18	63.97	66.19	102.84	141.35	141.12
$v_{\text{blackout}} = 100\,000$	53.54	71.47	34.89	51.82	146.56	65.29	67.26	104.29	143.84	143.40
$p_{\text{cap}} = 2\,500$	53.54	71.20	34.69	51.71	144.69	64.44	66.58	103.32	142.26	141.95
$p_{\text{cap}} = 10\,000$	53.54	71.32	34.77	51.79	145.20	64.91	66.93	103.90	143.08	142.70
$\zeta_j = 1$	53.54	70.17	33.64	51.54	157.45	70.04	72.47	113.94	155.14	154.52
c_{ref} = observed load	53.14	70.68	34.40	51.36	145.09	63.94	66.00	103.27	142.12	141.79
Two-sided price correction	52.15	69.84	34.24	50.34	144.43	63.93	66.05	104.49	144.97	144.59

Table 26: Consumption, shadow price, system cost, welfare-adjusted cost, and LCOE summary for the curtailment-cost, renewal-cost, and reference-consumption cases (baseline technology costs, 5% discount rate, $|\varepsilon| = 0.05$).

The price cap shifts welfare losses between the curtailment branch and the range of flexible demand (Table 27). With a price cap of 2 500 €/MWh, consumption lies on the curtailment branch in 59 hours per year on average, and the welfare loss there is 1.85 B€/y; with a price cap of 10 000 €/MWh, it lies there in 8 hours per year, with a loss of 1.20 B€/y, compared with 23 hours and 1.50 B€/y in the baseline. Across the price caps, the total welfare loss varies only between 3.33 and 3.44 B€/y. The price cap therefore mainly determines which consumption reductions count as curtailment and has little effect on their total welfare cost or on the optimal mix.

Case	Welfare loss $c \le k(\vartheta)$		Welfare loss $k(\vartheta) < c < c_{\text{ref}}(\vartheta)$		Total welfare loss		Welfare gain $c > c_{\text{ref}}(\vartheta)$	
	Annual [B€/y]	$/(T_Y \bar{c}_{\text{ref}})$ [€/MWh]	Annual [B€/y]	$/(T_Y \bar{c}_{\text{ref}})$ [€/MWh]	Annual [B€/y]	$/(T_Y \bar{c}_{\text{ref}})$ [€/MWh]	Annual [B€/y]	$/(T_Y \bar{c}_{\text{ref}})$ [€/MWh]
Baseline	1.50	3.20	1.90	4.05	3.40	7.25	1.30	2.78
$v_{\text{blackout}} = 25\,000$	1.55	3.30	1.96	4.19	3.51	7.49	1.29	2.75
$v_{\text{blackout}} = 100\,000$	1.45	3.10	1.83	3.89	3.28	6.99	1.32	2.81
$p_{\text{cap}} = 2\,500$	1.85	3.94	1.59	3.39	3.44	7.33	1.30	2.77
$p_{\text{cap}} = 10\,000$	1.20	2.56	2.12	4.53	3.33	7.09	1.31	2.79
$\zeta_j = 1$	1.60	3.41	2.08	4.43	3.68	7.84	1.24	2.65
c_{ref} = observed load	1.47	3.16	1.88	4.05	3.35	7.20	1.29	2.78
Two-sided price correction	1.50	3.28	1.88	4.11	3.37	7.39	1.26	2.75

Table 27: Monetary welfare losses and gains relative to reference consumption for the curtailment-cost, renewal-cost, and reference-consumption cases (baseline technology costs, 5% discount rate, $|\varepsilon| = 0.05$).

Each case evaluates welfare with its own consumption-benefit function and curtailment threshold. Paying the full initial CAPEX at each renewal raises annualized system cost from 64.59 to 70.04 B€/y and the welfare-adjusted LCOE from 142.19 to 154.52 €/MWh, an increase of 9%. The higher renewal cost weighs most on batteries, which have the shortest lifetime, and battery capacity falls from 44 to 38 GWh. Wind capacity falls from 312 to 305 GW and solar capacity rises from

216 to 223 GW, while hydrogen storage capacity, with a lifetime of 100 years, rises from 35.0 to 36.2 TWh.

In all seven cases, the optimal mix keeps the structure of the baseline: 305 to 313 GW of wind, 207 to 223 GW of solar, between 33 and 37 TWh of hydrogen storage, and a small battery capacity.

M Perfect-Foresight Optima

M.1 Perfect-Foresight Model

This appendix develops a perfect-foresight benchmark for the fully renewable system with hydrogen and battery storage. For a weather path $\mathcal{P}$ of consecutive years, observed or simulated, let $\hat{\psi}_{Wt}$ and $\hat{\psi}_{St}$ denote the periodic piecewise-linear interpolations of its hourly wind and solar capacity factors, with the value of each hour placed at its midpoint, and let $T_{\mathcal{P}} = T_Y|\mathcal{P}|$ hours. Renewable production along the known path is

$$Y_t^{\mathrm{PF}}(q) = K_W\hat{\psi}_{Wt} + K_S\hat{\psi}_{St}, \qquad 0 \le t \le T_{\mathcal{P}}. \tag{70}$$

The capacity vector is restricted as in the stochastic problem without fossil backup; see Appendix F.

For fixed capacities, the complete storage-dispatch path is chosen under complete information about every future realization in $\mathcal{P}$. Storage dynamics, bounds, and energy balance are those of Section 2, with $Y_t^{\mathrm{PF}}(q)$ in place of Y_t, and the stocks satisfy the cyclic conditions

$$S_{i0} = S_{iT_{\mathcal{P}}}, \qquad i \in \{H, B\}. \tag{71}$$

The cyclic condition prevents end-of-path storage depletion from artificially raising operating value. The optimized period-average hourly operating value is

$$\rho_{\mathcal{P}}^{q} = \sup_{a} \frac{1}{T_{\mathcal{P}}} \int_0^{T_{\mathcal{P}}} U_{\theta_t}(c_t)dt. \tag{72}$$

The capacity problem annualizes this period average and applies the same lifecycle capacity costs as the stochastic problem:

$$q_{\mathcal{P}}^{*} \in \arg\max_{q} \left\{ \frac{T_Y \rho_{\mathcal{P}}^{q}}{\delta} - C(q) \right\}. \tag{73}$$

Thus a five- or ten-year path affects the optimum through its average hourly operating value.

Four consecutive 15-minute observations are averaged to each UTC hour, and February 29 is omitted, giving 8 760 observations per year. The analysis covers five individual years, two five-year paths, and the full 2016-2025 path. Each path is repeated cyclically in the operating problem. In every weather year, reference consumption is that of the stochastic problem, derived from the 2025 hourly load (Section 3), so the paths differ only in renewable production. Demand elasticity is fixed at $|\varepsilon| = 0.05$.

M.2 Perfect-Foresight Results

Capacity cost is the annualized lifecycle cost. NPV and the three levelized cost measures are calculated by simulating the optimal operation of each selected optimum along its weather path, separately from the capacity optimization. Table 18 and Table 28 report the optima with baseline costs at discount rates of 5% and 2%, and Table 29 reports them with low battery costs at a 5% discount rate.

Period	Wind [GW]	Solar [GW]	H_2 storage [GWh]	Battery [GWh]	Electrolyzer [GW]	H_2 CCGT [GW]	Capacity cost [B€/y]	NPV [T€]	LCOE [€/MWh]	$LCOE_{CE}$ [€/MWh]	$LCOE_{adj}$ [€/MWh]
2017	309	212	21 928	32	49	50	41.42	−2.07	70.55	88.07	88.06
2019	275	213	16 047	50	39	47	37.27	−1.85	61.98	78.74	78.69
2021	341	195	18 762	35	46	42	41.83	−2.09	69.57	88.91	88.91
2023	263	207	18 364	29	46	48	36.90	−1.82	64.41	77.86	77.80
2025	306	235	23 172	42	45	47	41.57	−2.07	73.77	88.29	88.29
2016-20	325	231	27 576	35	45	46	43.27	−2.17	66.78	92.33	92.33
2021-25	305	216	23 604	29	48	48	41.10	−2.05	69.07	87.42	87.41
2016-25	318	225	26 006	32	46	47	42.52	−2.12	67.64	90.58	90.58

Table 28: Perfect-foresight optima with baseline costs and a 2% discount rate.

Period	Wind [GW]	Solar [GW]	H_2 storage [GWh]	Battery [GWh]	Electrolyzer [GW]	H_2 CCGT [GW]	Capacity cost [B€/y]	NPV [T€]	LCOE [€/MWh]	$LCOE_{CE}$ [€/MWh]	$LCOE_{adj}$ [€/MWh]
2017	276	195	23 909	287	41	40	56.92	−1.15	104.04	123.04	122.81
2019	248	186	16 543	363	28	36	50.03	−1.01	89.31	107.42	107.36
2021	314	170	17 229	434	31	29	56.85	−1.15	99.01	122.94	122.72
2023	224	220	19 109	305	35	40	50.38	−1.01	93.79	108.10	108.04
2025	276	229	22 847	342	32	35	56.78	−1.15	104.36	122.68	122.46
2016-20	280	238	33 126	344	28	35	59.74	−1.21	96.86	129.75	129.43
2021-25	276	204	24 563	291	38	39	56.88	−1.15	100.39	123.14	122.91
2016-25	279	225	30 374	316	32	38	58.87	−1.20	98.82	127.85	127.55

Table 29: Perfect-foresight optima with low battery costs and a 5% discount rate.

Lower battery costs move optimal battery capacity from tens to hundreds of GWh, while hydrogen storage remains at the much larger scale of tens of TWh. In both cases, the perfect-foresight optimum for 2016-2025 holds markedly less hydrogen storage than the corresponding optimum under uncertainty (26 006 rather than 38 909 GWh at a 2% discount rate, and 30 374 rather than 40 047 GWh with low battery costs), as in the baseline at 5% in Section 6.

N Observing the Drought Regime

The operator in Section 2 observes the drought regime. A drought is not announced, however: low wind and solar production may reflect an ordinary weather fluctuation or the onset of a persistent drought. This appendix compares the results with those of an operator who observes the phase, the storage levels, and the latent state X_t, but not the regime, and who assesses the regime from the current weather state through the belief

$$\varpi(x,\vartheta)=\frac{\bar{\varpi}e^{-\|x+m(\vartheta)\|^2/2}}{(1-\bar{\varpi})e^{-\|x\|^2/2}+\bar{\varpi}e^{-\|x+m(\vartheta)\|^2/2}}, \tag{74}$$

where $\bar{\varpi}=\gamma_{01}/(\gamma_{01}+\gamma_{10})$ is the stationary share of the scarcity regime. The belief treats the latent state in each regime as normally distributed around the regime's target, 0 or $-m(\vartheta)$, with identity covariance. It uses only the current weather state and ignores the information contained in past weather and in the storage levels.

The belief-based policy takes the same action in both regimes. It replaces the regime-specific shadow values of stored energy in (10) of Section 2 by their belief-weighted average:

$$\pi^q_{\mathrm{b}}(x,s,\vartheta)\in\arg\max_{(\alpha,\nu)\in A(\xi)}\left\{U_\vartheta(y(\xi)+\nu-\alpha_H-\alpha_B)-p_f\nu+\sum_{i\in\{H,B\}}b_i(\alpha_i)\bar{\lambda}_i\right\}, \tag{75}$$

$$\bar{\lambda}_i=(1-\varpi(x,\vartheta))\partial_{s_i}W^q_{\mathrm{b}}(x,0,s,\vartheta)+\varpi(x,\vartheta)\partial_{s_i}W^q_{\mathrm{b}}(x,1,s,\vartheta).$$

Production and the admissible actions do not depend on the regime, so the policy can be implemented without knowledge of the regime. Here W^q_{b} is the relative value of the belief-based policy in each regime: it solves the equation obtained from (10) in Section 2 by evaluating the Hamiltonian at π^q_{b} instead of taking the supremum, and ρ^q_{b} denotes the long-run payoff of π^q_{b}. The policy and its relative value are determined jointly and computed with the scheme of Appendix F, with the belief tabulated at the mid-hour phase of every hour.

The belief-based policy is optimal given the belief at each instant. Because the belief is not the conditional probability of the regime given the operator's information, the belief-based policy need not be the best one based on that information. It is also feasible for an operator who observes the regime, so $\rho^q_{\mathrm{b}}\le\rho^q$ at given capacities.

The capacities are optimized separately for each operator. Under the belief-based policy, the capacity gradient of Appendix F is incomplete: the common action maximizes the belief-weighted objective rather than the objective of each regime, so its response to the capacities does not cancel in the derivative of ρ^q_{b}. The gradient-based optimization of capacities for this operator therefore uses the exact gradient, which includes this response.

The comparison uses the system with battery and hydrogen storage and without fossil backup, at baseline costs and a 5% discount rate. Table 30 reports the optimal capacities, the NPV, the LCOE on certainty-equivalent consumption, and the welfare-adjusted LCOE for both operators.

Operator	Wind [GW]	Solar [GW]	H_2 storage [GWh]	Battery [GWh]	Electrolyzer [GW]	H_2 CCGT [GW]	NPV [T€]	$\mathrm{LCOE}_{\mathrm{CE}}$ [€/MWh]	$\mathrm{LCOE}_{\mathrm{adj}}$ [€/MWh]
Regime observed	312	216	35 010	44	47	45	−1.334	142.52	142.19
Belief	313	216	35 211	43	47	46	−1.335	142.68	142.34

Table 30: Optimal capacities, NPV, LCOE on certainty-equivalent consumption, and welfare-adjusted LCOE with and without observation of the drought regime: battery and hydrogen storage without fossil backup, baseline costs, 5% discount rate.

Observing the regime barely changes the optimal mix. Both operators choose nearly the same wind, solar, electrolyzer, and hydrogen CCGT capacities, which differ by at most 1 GW; the belief-based operator holds slightly more hydrogen storage and slightly less battery capacity. As a result, the NPV is only slightly higher when the regime is observed. This gap is an upper bound on the welfare cost of not observing the regime, and the capacities chosen for the belief-based policy need not be optimal for an operator who uses the available information optimally.

O Lifetime Cost of Fossil Electricity

This appendix computes the lifetime cost of meeting reference consumption with natural gas, hard coal, or lignite alone. Capacity equals peak reference consumption, so the plant fleet runs $T_Y \bar{c}_{\text{ref}} / \max_\vartheta c_{\text{ref}}(\vartheta) = 6\,013$ full-load hours per year, with the mean of 53.5 GW and the peak of 78.0 GW of Section 3. With the lifecycle capacity cost C_j of (47) in Appendix D, $\delta = 0.05$, and $\zeta_j = 0.5$, the cost per MWh of technology j is

$$\text{LCOE}_j = \frac{\delta C_j \max_\vartheta c_{\text{ref}}(\vartheta)}{T_Y \bar{c}_{\text{ref}}} + \frac{p_j}{\eta_j} + v_j + \tau\gamma_j, \tag{76}$$

where p_j is the fuel price per MWh of fuel, η_j the net electrical efficiency, v_j variable O&M, and γ_j the emissions per MWh of electricity. For natural gas, γ_j is the intensity γ_f of Section 7.

Plant data for natural-gas CCGT are those of Section 7 (Danish Energy Agency, 2026a); for hard coal and lignite, they are the midpoints of the 2024 CAPEX ranges and the remaining 2024 parameters of Kost et al. (2024). The natural-gas prices are the EU import price in USD/MMBtu (International Monetary Fund, 2026), converted at the monthly exchange rate (Board of Governors of the Federal Reserve System, 2026), for July 2026 and averaged over 2023 to 2025. All prices are in 2025 euros.[41] The hard-coal price is the average import price in the first half of 2026 (Statistisches Bundesamt, no date); lignite is not traded, and its price is taken from Kost et al. (2024). Emissions of hard coal and lignite follow from the 2020 fuel emission factors for raw hard coal in power plants (93.6 tCO_2/TJ) and for Rhineland raw lignite in public power plants (113.3 tCO_2/TJ) (Juhrich, 2022). The carbon price is the volume-weighted average German auction price of August 2026, 82.31 €/tCO_2 (Umweltbundesamt, 2026), or 79.38 €/tCO_2 in 2025 euros.

Technology	CAPEX [€/kW]	Fixed O&M [€/(kW·y)]	Lifetime [y]	η_j	v_j [€/MWh]	p_j [€/MWh_{th}]	γ_j [tCO_2/MWh]
Natural-gas CCGT	1 092	36.4	25	56%	5.5	57.5 / 42.1	0.36
Hard coal	2 042	37.8	30	39%	5.1	16.9	0.86
Lignite	2 247	42.9	40	38%	5.1	2.3	1.07

Table 31: Plant and fuel data for fossil-only supply. The two natural-gas prices are those of July 2026 and the 2023 to 2025 average. Fuel prices p_j are per MWh of fuel energy (MWh_{th}). Emissions γ_j are per MWh of electricity.

[41] The natural-gas CCGT data are converted as in Section 3, and all other prices with the euro-area HICP (Eurostat, 2026a; 2026b): the data of Kost et al. (2024) from 2024, the 2023-2025 gas prices with the annual index of their year, the July 2026 gas price with the index of that month, the hard-coal price with the average index of January to June 2026, and the carbon price with the index of August 2026. Both gas prices are also converted from the gross to the net calorific value of natural gas with the factor 1.108 of Juhrich (2022), as in Section 7.

Technology	Capacity	Fuel	Variable O&M	CO_2 at 79.38 €/t	Total	
					$\tau = 0$	$\tau = 79.38$
	[€/MWh]	[€/MWh]	[€/MWh]	[€/MWh]	[€/MWh]	[€/MWh]
Natural gas, July 2026	16.9	102.6	5.5	28.6	125.0	153.6
Natural gas, 2023-2025	16.9	75.1	5.5	28.6	97.5	126.1
Hard coal	25.7	43.4	5.1	68.6	74.2	142.8
Lignite	27.3	6.2	5.1	85.2	38.6	123.8

Table 32: Lifetime cost of meeting reference consumption with a single fossil fuel (5% discount rate, 6 013 full-load hours), without a carbon price and at the August 2026 carbon price. Components may not add up to the totals because of rounding.

The welfare-adjusted LCOE of (21) in Section 2 follows from optimizing capacity against demand. With marginal cost $m_j = p_j/\eta_j + v_j + \tau\gamma_j$ and fossil capacity F, the planner solves

$$\max_{F \geq 0} \frac{1}{\delta} \int_0^{T_Y} \left[U_\vartheta(c_F(\vartheta)) - m_j c_F(\vartheta)\right] d\vartheta - C_j F, \qquad c_F(\vartheta) = \min\left\{F, c_{\text{ref}}(\vartheta)\left(p_{\text{ref}}/m_j\right)^{|\varepsilon|}\right\}, \tag{77}$$

where $c_F(\vartheta)$ maximizes $U_\vartheta(c) - m_j c$ subject to $c \leq F$; the closed form holds because $m_j < p_{\text{cap}}$. The optimal capacity F_j^* satisfies $\int_0^{T_Y} \max\left\{U'_\vartheta\left(F_j^*\right) - m_j, 0\right\} d\vartheta = \delta C_j$, and

$$\text{LCOE}_{\text{adj},j} = \frac{\delta C_j F_j^* - \int_0^{T_Y} \left[U_\vartheta\left(c_{F_j^*}(\vartheta)\right) - m_j c_{F_j^*}(\vartheta)\right] d\vartheta}{T_Y \bar{c}_{\text{ref}}}. \tag{78}$$

In Table 33, the optimal capacity of 68 to 69 GW lies below the peak reference consumption of 78.0 GW. It binds in about 6% to 19% of hours, when consumption falls short of its unconstrained level. Because capacity is chosen optimally and consumption responds to the marginal cost, the welfare-adjusted LCOE lies about 1.5 to 5 €/MWh below the lifetime cost in Table 32.

Technology	$\tau = 0$		$\tau = 79.38$	
	F_j^* [GW]	$\text{LCOE}_{\text{adj},j}$ [€/MWh]	F_j^* [GW]	$\text{LCOE}_{\text{adj},j}$ [€/MWh]
Natural gas, July 2026	68.7	123.4	68.4	151.5
Natural gas, 2023-2025	69.1	96.0	68.7	124.4
Hard coal	68.6	71.2	67.7	139.8
Lignite	69.2	33.5	67.8	120.9

Table 33: Welfare-optimal capacity and welfare-adjusted LCOE of fossil-only supply, without a carbon price and at the August 2026 carbon price (5% discount rate, inputs of Table 31).

For comparison, the hourly German day-ahead prices of Section 3, which apply to electricity from all generation technologies, average 82.3 €/MWh over 2015-2025, 89.7 €/MWh over 2023-2025, and 89.3 €/MWh in 2025 (European Network of Transmission System Operators for Electricity (ENTSO-E), 2026).[42]

[42]Prices of earlier years are converted to 2025 euros with the annual euro-area HICP (Eurostat, 2026a).