



ARTICLE

Health-Aware Dynamic Operation of PEM Electrolyzers: A Techno-Environmental Optimization for Carbon-Matched Hydrogen Production

Cosimo Magazzino^{1,2,3,*}, Muhammad Waqas^{4,5} and Mustafa Tevfik Kartal⁶

¹Department of Management, Finance and Technology, LUM University “Giuseppe Degennaro”, Casamassima, Italy

²Economic Research Center, Western Caspian University, Baku, Azerbaijan

³ARUCAD Research Centre, Arkin University of Creative Arts and Design, Kyrenia, Northern Cyprus, Türkiye

⁴Department of Environmental Sciences, Kohat University of Science and Technology, Kohat, Pakistan

⁵Department of Urban and Regional Planning, Karadeniz Technical University, Trabzon, Türkiye

⁶Department of Finance and Banking, European University of Lefke, Lefke, Northern Cyprus, Türkiye

*Corresponding Author: Cosimo Magazzino. Email: magazzino@lum.it

Received: 21 May 2026; Accepted: 29 July 2026; Published: 18 October 2026

ABSTRACT: Hydrogen is increasingly positioned as a strategic energy carrier for decarbonizing hard-to-abate sectors, yet the actual performance of electrolytic hydrogen depends on how electrolyzers interact with renewable electricity, carbon accounting rules, storage limits, and stack degradation. This study develops a calibrated health-aware carbon-matched dispatch framework for proton exchange membrane (PEM) electrolysis. The model integrates hourly renewable generation, grid carbon intensity, part-load operation, start-stop degradation, hydrogen storage, and system cost in a single techno-environmental optimization structure. A reference 10 MW PEM electrolyzer coupled with solar and wind generation is simulated over 8760 hourly periods and assessed under four operating strategies: baseload grid-assisted electrolysis, price-only flexible operation, carbon-aware flexible operation, and the proposed health-aware carbon-matched strategy. To strengthen robustness, the model is calibrated against external PEM electrolysis techno-economic and operational benchmarks and then stress-tested under conservative efficiency, low-degradation, and strict-carbon scenarios. Results show a sharp trade-off. Baseload operation yields the lowest apparent levelized cost of hydrogen, 4.58 USD kg H₂⁻¹, but produces high operational emissions, 4.887 kg CO₂ kg H₂⁻¹. Carbon-aware flexible operation cuts emissions, but suffers from low utilization and frequent shutdowns. The proposed health-aware carbon-matched strategy reduces operational emissions to 0.027 kg CO₂ kg H₂⁻¹, maintains 99.8% direct use of renewable electricity, and lowers the levelized cost of hydrogen (LCOH) by about 31% relative to the carbon-aware flexible benchmark. The central finding is direct: flexibility alone is not enough. Renewable hydrogen systems must be operated to respect both hourly carbon conditions and stack health. Within the operational electricity-related boundary, this produces hydrogen with lower reported emissions, fewer damaging operating events, and stronger techno-economic credibility. The reported emissions are not cradle-to-grave lifecycle emissions; conclusions about full environmental superiority require a dedicated lifecycle assessment.

KEYWORDS: Green hydrogen; PEM electrolysis; dynamic operation; stack degradation; hourly carbon matching; renewable hydrogen; techno-environmental optimization; hydrogen storage

1 Introduction

Hydrogen has become one of the most contested vectors in the energy transition. It is not a universal solution, and it should not be treated as one. Yet some decarbonization problems remain hard to solve with direct electrification alone. Ammonia, methanol, refining, direct-reduced iron, long-duration storage, synthetic fuels, shipping, and parts of the chemical sector need molecules, process heat, feedstocks, or dispatchable energy carriers [1,2]. In those settings, hydrogen can do something electricity cannot easily do on its own: it can connect renewable electricity to material and fuel systems.

That conditionality is important. Hydrogen is not clean by default. Its climate value depends on the production route, the electricity source, the timing of electricity consumption, storage configuration, transport losses, water sourcing, and final use. Global hydrogen production is still dominated by fossil-based routes. The International Energy Agency (IEA) reports that low-emissions hydrogen production grew in 2024 and was expected to approach 1 Mt in 2025, but it still accounts for less than 1% of total hydrogen production [3]. The gap between policy ambition and the operation of low-carbon assets is therefore not a communication problem. It is a system-design problem.

Water electrolysis is central to this design problem. When powered by additional renewable electricity or very-low-carbon grids, electrolysis can supply hydrogen with much lower operational emissions than unabated fossil-fuel routes. Among electrolyzer technologies, proton exchange membrane (PEM) electrolysis is particularly relevant for renewable coupling because it can respond quickly, operate at high current densities, and tolerate dynamic power inputs better than many conventional alkaline configurations [4,5]. This fast response is valuable. It is also dangerous when interpreted too casually. A stack that can ramp quickly is not a stack that should be forced into unlimited start-stop operation.

The paper addresses this operational tension. Current green hydrogen models often optimize capacity factor, electricity cost, renewable matching, or carbon intensity [6,7]. These metrics are necessary, but they do not fully describe the problem. A high capacity factor can lower the levelized cost of hydrogen (LCOH) while raising emissions if grid electricity is carbon-intensive. A strict carbon rule can reduce emissions, but lower utilization and increase cycling. A price-responsive strategy can look rational in an electricity market while quietly shortening stack life. The missing layer is health-aware dispatch: control logic that uses flexibility without damaging the electrochemical asset that enables it.

Recent hydrogen research has increasingly moved beyond purely technical cost analysis. Horky et al. [8] show that financial market development can affect hydrogen adoption, indicating that deployment depends not only on technology readiness but also on capital market conditions and investment channels. Complementing this perspective, Magazzino and Haroon [9] discuss how AI-based modeling, machine learning, and digital twin approaches can accelerate hydrogen process optimization and improve the discovery and control of production technologies. At the territorial scale, Magazzino et al. [10] document substantial regional disparities in hydrogen production capacity across 27 European countries, highlighting the importance of infrastructure, industrial structure, and policy support. Together, these studies reinforce the need for hydrogen-system models that integrate technological performance, economic feasibility, and spatial or institutional heterogeneity.

The research question is, therefore, precise: how should a renewable PEM electrolyzer be operated to produce hydrogen that is simultaneously low-carbon, stack-conscious, and cost-credible? The answer proposed here is a calibrated health-aware carbon-matched optimization framework. It allows the electrolyzer to run during renewable-rich, low-carbon windows, but it discourages isolated short operating events and penalizes starts and sharp ramps. Hydrogen storage is included because it alters the plant's temporal structure. It allows production to be consolidated into cleaner, more stable operating blocks, rather than forcing the electrolyzer to follow demand or renewable fluctuations on an hour-by-hour basis.

The novelty of this study is threefold. First, unlike previous techno-economic studies that primarily optimize hydrogen cost, utilization, or renewable electricity supply, this paper directly embeds PEM stack health exposure into the dispatch problem. Starts, ramps, and operating exposure are therefore treated as endogenous determinants of system performance, not as ex post discussion variables. Second, the study extends carbon-aware electrolysis dispatch by shifting the operational criterion from isolated clean-hour selection to storage-enabled clean-block scheduling. This distinction matters because a dispatch rule can be low-carbon but still operationally damaging if it fragments production into many short operating periods. Third, the framework links hourly carbon accounting, storage operation, degradation-sensitive control, and LCOH within a calibrated robustness architecture. This integrated treatment allows the paper to identify a system-level trade-off that previous studies have not fully quantified: renewable hydrogen should be simultaneously carbon-matched, stack-conscious, and storage-enabled to be operationally low-carbon and economically durable. Full environmental credibility, however, also requires cradle-to-grave lifecycle assessment.

The rest of the paper is organized as follows. [Section 2](#) reviews the literature on renewable hydrogen, temporal carbon accounting, dynamic PEM electrolysis, and environmental assessment. [Section 3](#) develops the conceptual framework underlying the health-aware carbon-matched dispatch strategy. [Section 4](#) presents the reference system, input data, dispatch strategies, formal optimization model, emissions boundary, LCOH parameterization, calibration, and degradation assumptions. [Section 5](#) reports the main techno-environmental results, storage-related mechanisms, and benchmarking against the recent PEM electrolysis literature. [Section 6](#) develops the sensitivity and robustness analysis, supported by additional three-dimensional visualizations of the strategy frontier, storage effects, sensitivity landscape, and robustness scenarios. [Section 7](#) discusses the scientific, technological, policy, water, storage, and lifecycle implications of the findings. [Section 8](#) identifies the main limitations and future research directions. [Section 9](#) concludes.

2 Literature Background and Research Gap

2.1 Renewable Hydrogen and Temporal Carbon Accounting

Early green hydrogen assessments often asked whether electrolysis could become cost-competitive. That question remains relevant, but it is no longer enough. Electrolytic hydrogen is embedded in power systems whose marginal emissions vary by hour. A hydrogen plant that buys electricity from the grid during fossil-dominated hours can still have high operational emissions, even if it also purchases renewable certificates annually. The shift from annual averages to temporal accounting is therefore one of the most important methodological changes in hydrogen research.

Temporal matching research shows why. Annual matching can mask large intra-year mismatches between hydrogen production and renewable generation. Monthly, daily, and hourly matching progressively reduce this problem but also increase operational complexity. Casas Ferrus et al. [7] show that hourly matching need not be prohibitively costly when renewable portfolios are diversified across technologies and locations. Their result is highly relevant here because it implies that strict carbon accounting should not be analyzed as a single constraint in isolation. It interacts with the renewable mix, storage, and operating strategy.

There is also a project-delivery dimension. Odenweller and Ueckerdt [11] document a substantial gap between announced and completed green hydrogen capacity. The problem is usually described in terms of cost, finance, and demand uncertainty. Those factors are real. Yet operational design also matters. If projects are assessed using annual averages and generic capacity-factor assumptions, they may appear bankable before the hourly carbon, storage, and degradation issues are fully visible.

2.2 PEM Electrolysis under Dynamic Operation

PEM electrolysis has several attributes that suit renewable coupling: high current density, compact design, high-purity hydrogen output, rapid response, and differential-pressure operation. The DOE Hydrogen Shot assessment identifies water electrolysis as a central pathway for clean hydrogen production and emphasizes the need for improvements in cost, efficiency, durability, and manufacturing [4]. These improvements are connected. A cheaper electrolyzer that degrades quickly under real operation may not lower hydrogen cost in practice.

Dynamic operation is the critical interface. PEM stacks can respond rapidly to changing power input, but variable operation can impose stress through start-stop events, voltage cycling, pressure changes, thermal gradients, membrane hydration transients, gas crossover, and catalyst degradation [12]. Sayed-Ahmed et al. [5] review the literature and emphasize the need to link dynamic operating conditions to equipment behavior, degradation, and electricity market participation. The implication is blunt: dispatch is not only an economic choice. It is also a durability choice.

This creates a gap in many techno-economic models. They may include stack replacement after a fixed lifetime or assume a constant degradation rate. Such assumptions are simple, but they erase the difference between stable operation and fragmented operation. Two plants with the same annual capacity factor can have very different operating histories. If one plant runs in long blocks and the other starts hundreds of times, a fixed-lifetime assumption misrepresents their relative economics and environmental burden.

2.3 Environmental Assessment and Lifecycle Credibility

Hydrogen environmental assessment has also become more granular. Wei et al. [13] assess future lifecycle impacts of global hydrogen production across regions and pathways and show that electrolysis becomes progressively less carbon-intensive as electricity systems decarbonize. This is not an automatic result; it depends on the electricity system. Fossil-based hydrogen without carbon capture has limited emissions-reduction potential, while electrolysis improves as grids and renewable supplies become cleaner.

Operational emissions from electricity consumption are only one layer of the problem. Equipment manufacturing, stack replacement, water purification, compression, storage, transport, and downstream conversion also matter. For a study focused on operations, it is reasonable to model hourly operational emissions directly and address lifecycle effects through sensitivity analyses and discussion. It would be misleading, however, to call operational emissions full lifecycle emissions.

Water is another often-compressed issue. Stoichiometry requires about 9 kg of water per kg of hydrogen, but real systems need additional water for purification, cooling, and losses. In regions with excellent solar resources and water stress, water sourcing can become an environmental constraint. The framework developed here includes water cost but treats spatial water stress as a limitation and a future extension.

2.4 Research Gap

The specific gap addressed in this paper is the absence of an integrated dispatch framework that simultaneously accounts for hourly carbon intensity, PEM stack-health exposure, hydrogen storage, and hydrogen cost. Prior studies have advanced important insights on renewable hydrogen economics, temporal matching, PEM degradation, and environmental assessment, but these elements are often modeled separately. This paper contributes by embedding degradation-sensitive operation directly into a carbon-matched dispatch model and by showing that clean-hour selection alone is insufficient. The proposed health-aware strategy organizes hydrogen production into stable, low-carbon blocks, thereby reducing emissions, start-up degradation exposure, and costs relative to a purely carbon-aware, flexible benchmark.

3 Conceptual Framework

The conceptual framework is built around the idea that renewable hydrogen is not produced solely by an electrolyzer. It is the output of a coupled energy system. The electrolyzer, renewable generator, grid connection, storage vessel, demand profile, and control algorithm jointly determine cost, emissions, and durability.

Four propositions guide the model. First, hydrogen should not be classified as low-carbon solely on the basis of annual-average electricity assumptions. Temporal conditions matter because electricity-system carbon intensity is not constant. Second, a PEM electrolyzer is not a passive flexible load. Its operating history affects efficiency, degradation, and replacement cost. Third, hydrogen storage is not an auxiliary component; it changes the feasible set by decoupling production from demand. Fourth, no single metric is sufficient. LCOH, carbon intensity, utilization, renewable share, and degradation exposure must be read together.

The proposed strategy, health-aware carbon-matched dispatch, implements these propositions. It searches for low-carbon production blocks rather than isolated low-carbon hours. It avoids ultra-low-load operation. It limits unnecessary starts. It uses storage to buffer demand and renewable variability. It also enables a transparent comparison with simpler rules common in hydrogen analysis, such as baseload utilization, price arbitrage, and carbon-only filtering.

The conceptual claim is not that the proposed strategy is universally optimal across all geographies. It is the dispatch strategy that internalizes stack-health exposure that can change the ranking of hydrogen-system designs. That is a scientific claim because it is testable under alternative parameter values, and it is a policy-relevant claim because subsidy and certification rules can unintentionally reward operating profiles that appear clean or cheap but are not durable. Fig. 1 gives a graphical description of the theoretical background.




Figure 1: Conceptual framework: health-aware carbon-matched PEM electrolysis.

4 Materials and Methods

4.1 Reference System and Data Structure

Input-Data Construction

The simulation uses an hourly input matrix with 8760 observations, corresponding to one full operating year. The dataset is designed as a calibrated reference profile rather than as a site-specific measured dataset. This choice allows the four dispatch strategies to be compared under identical physical and operational conditions. Each hourly observation contains renewable generation availability, electricity price, grid carbon intensity, hydrogen demand, and water availability. These variables are treated as exogenous inputs in the optimization.

Renewable generation is represented by separate profiles for photovoltaic and wind. Available renewable electricity in hours t is calculated as the sum of PV and wind output:

$$R_t = K_{PV} CF_{PV,t} + K_W CF_{W,t}$$

where R_t is available renewable generation in MWh, K_{PV} and K_W are installed PV and wind capacities, and $CF_{PV,t}$ and $CF_{W,t}$ are hourly capacity factors. In the reference case, $K_{PV} = 18$ and $K_W = 12$ MW. The PV profile follows a daytime production pattern with seasonal amplitude, while the wind profile is allowed to vary more irregularly across hours. The use of two renewable technologies is intentional: it avoids making the electrolyzer dependent on a single-resource profile and reflects the portfolio logic of renewable hydrogen systems.

Electricity prices are included as an hourly exogenous variable p_t , expressed in USD MWh⁻¹. The price series affects the objective function whenever grid imports are permitted. It is especially relevant for the price-only flexible strategy, where the electrolyzer responds to low-price periods, and for the baseload grid-assisted strategy, where the grid supplies electricity whenever dedicated renewable generation is insufficient. In the carbon-aware and health-aware strategies, price remains part of the objective but is subordinated to carbon admissibility and stack-health constraints.

Grid carbon intensity is represented by an hourly emissions factor CI_t , expressed in kg CO₂ MWh⁻¹. This variable serves two roles. First, it determines the operational emissions associated with grid imports:

$$CO_{2,t} = P_{grid,t} CI_t$$

Second, it defines the carbon-admissibility rule used in the carbon-aware and health-aware dispatch cases. In the reference case, grid electricity is admissible only when hourly carbon intensity is below the calibrated threshold of 80 kg CO₂ MWh⁻¹. A stricter threshold of 50 kg CO₂ MWh⁻¹ is tested in the robustness analysis. This structure allows the model to distinguish between generic grid use and low-carbon grid support.

Hydrogen demand is specified as an hourly demand profile q_t , expressed in kg H₂ h⁻¹. Demand can be satisfied by contemporaneous electrolyzer production or by discharge from hydrogen storage. If neither production nor stored hydrogen is sufficient, the model uses an unmet-demand slack variable U_t , which is heavily penalized in the objective function. The inclusion of U_t prevents infeasibility and makes reliability trade-offs explicit. In the baseline case, demand is treated as an industrial offtake profile rather than as fully flexible demand. This assumption is conservative because it requires the storage system to absorb part of the mismatch between renewable production and hydrogen use.

[Table 1](#) summarizes the main input variables used in the reference simulation.

Table 1: Input data used in the hourly dispatch simulation.

Input Variable	Symbol	Unit	Baseline Representation	Role in the Model
PV installed capacity	K_{PV}	MW	18 MW	Defines maximum solar contribution
Wind installed capacity	K_W	MW	12 MW	Defines maximum wind contribution
PV capacity factor	$CF_{PV,t}$	Dimensionless	Hourly normalized solar profile	Determines hourly PV output
Wind capacity factor	$CF_{W,t}$	Dimensionless	Hourly normalized wind profile	Determines hourly wind output
Available renewable generation	R_t	MWh h ⁻¹	$K_{PV}CF_{PV,t} + K_WCF_{W,t}$	Upper bound for renewable electricity use
Electricity price	p_t	USD MWh ⁻¹	Hourly exogenous profile	Determines the cost of grid imports
Grid carbon intensity	CI_t	kg CO ₂ MWh ⁻¹	Hourly exogenous profile	Determines emissions and carbon admissibility
Hydrogen demand	q_t	kg H ₂ h ⁻¹	Hourly industrial demand profile	Must be met by production, storage discharge, or unmet-demand slack
Water availability/cost	W_t, c_w	kg H ₂ O; USD kg H ₂ O ⁻¹	Proportional to hydrogen output	Accounts for the water requirement of electrolysis

Note: The table reports the exogenous input variables used in the 8760-h reference dispatch simulation. All operating strategies share the same assumptions about renewable generation, prices, carbon intensity, demand, and water inputs.

4.2 Dispatch Strategies

Four strategies are evaluated. Strategy A is baseload grid-assisted electrolysis. The electrolyzer operates at near-constant load and imports grid electricity whenever dedicated renewable output is insufficient. Strategy B is price-only flexible electrolysis. The plant runs during low-price periods and reduces or stops production during expensive periods. Strategy C is carbon-aware flexible electrolysis. The plant operates only when renewable availability and carbon intensity meet a predefined threshold. Strategy D is the proposed health-aware carbon-matched strategy. It operates in clean production blocks, respects a minimum stable load, penalizes excessive starts, and uses storage to avoid mechanically following every short renewable fluctuation.

Practical Implementation of the Health-Aware Carbon-Matched Strategy

The proposed health-aware carbon-matched strategy is intended as a practical supervisory dispatch rule, not only as an abstract optimization case. In an operating plant, it can be implemented within an energy management system or a model predictive control layer. At each scheduling update, the controller receives renewable-generation forecasts, electricity-price data, hourly grid-carbon-intensity signals, hydrogen-demand requirements, storage state of charge, and stack-condition indicators such as voltage,

pressure, temperature, and recent start-up history. It then converts these inputs into hourly electrolyzer set-points that satisfy carbon restrictions while avoiding unnecessary cycling.

The first practical step is carbon admissibility screening. Hours are classified as admissible when the grid carbon intensity condition is satisfied and renewable availability is sufficient to support stable operation. This screening can be represented as:

$$a_t = \mathbf{1} \{ CI_t \leq \overline{CI}, R_t \geq P_{\min} \} \quad (1)$$

The admissibility indicator is not used to mechanically switch the electrolyzer on and off. Instead, consecutive admissible hours are grouped into candidate low-carbon production blocks:

$$B_k = \{ t_k, t_k + 1, \dots, t_k + L_k - 1 : a_t = 1 \} \quad (2)$$

This block logic is the practical difference between the proposed strategy and a conventional carbon-aware rule. A simple carbon-aware controller may run whenever an individual hour is clean. The health-aware strategy asks whether that hour belongs to a sufficiently stable operating block. Very short isolated events are therefore discouraged because their hydrogen benefit is small relative to the start-up and ramping stress imposed on the stack.

During an accepted production block, the electrolyzer power set-point must remain between the calibrated minimum stable load and the rated capacity:

$$P_{\min} y_t \leq P_{\text{el},t} \leq P_{\max} y_t \quad (3)$$

The minimum-load condition has a practical interpretation. It prevents the dispatch algorithm from keeping the stack energized at unrealistically low power levels only to satisfy a mathematical carbon rule. If the unit is on, it must operate in a range that is technically meaningful for hydrogen production and stack management.

The health-aware component is introduced through the degradation-exposure term. The controller penalizes three operating features: baseline energized operation, start-up events, and load ramps. In compact form, the hourly degradation proxy is:

$$D_t = \alpha \frac{P_{\text{el},t}}{P_{\max}} + \beta z_t + \gamma r_t \quad (4)$$

where the start and ramp variables are linked to observable operating transitions:

$$z_t \geq y_t - y_{t-1}, r_t \geq \frac{P_{\text{el},t} - P_{\text{el},t-1}}{P_{\max}}, r_t \geq \frac{P_{\text{el},t-1} - P_{\text{el},t}}{P_{\max}} \quad (5)$$

In practice, the algorithm therefore follows a rolling sequence. First, it forecasts renewable availability, carbon intensity, demand, and storage over a day-ahead or intraday horizon. Second, it identifies admissible low-carbon windows. Third, it groups these windows into production blocks and rejects isolated starts unless they are needed for reliability. Fourth, it calculates electrolyzer setpoints that account for minimum load, storage capacity, demand, and degradation penalties. Fifth, the selected schedule is passed to the plant control system, while lower-level safety controllers enforce manufacturer limits on pressure, temperature, water flow, gas purity, and compressor operation.

This practical interpretation also clarifies the role of hydrogen storage. Storage does not merely increase capacity; it changes the operating logic. When storage has available capacity, the electrolyzer can produce

during a clean and stable block even if contemporaneous hydrogen demand is lower. When storage is depleted, the controller can prioritize production in the next admissible block rather than forcing the stack to start during a carbon-intensive or highly fragmented period. The result is a strategy that is carbon-matched, storage-enabled, and stack-conscious at the same time.

The strategy is not meant to replace plant-level protection systems. It operates above them. The optimization layer proposes hourly or sub-hourly set-points; the supervisory control and data acquisition system and balance-of-plant controllers enforce real-time safety constraints. If stack voltage, temperature, differential pressure, hydrogen purity, or compressor status moves outside acceptable limits, those lower-level systems override the dispatch schedule. In that sense, the proposed strategy is directly implementable as a supervisory operating policy while remaining compatible with conventional industrial control architecture.

4.3 Formal Optimization Model

The dispatch problem is formulated as a deterministic mixed-integer linear optimization model over hourly periods $t = 1, \dots, T$, with $T = 8760$. The model determines the hourly operating schedule of the PEM electrolyzer, the use of renewable and grid electricity, hydrogen production, hydrogen storage operation, renewable curtailment, and unmet hydrogen demand. The model is solved separately for each dispatch strategy to isolate the effect of the operating rule under the same reference-system configuration.

The main decision variables are electrolyzer power input $P_{el,t}$, renewable electricity used by the electrolyzer $P_{ren,t}$, grid electricity imported $P_{grid,t}$, curtailed renewable electricity $P_{curt,t}$, hydrogen production H_t , hydrogen storage level S_t , hydrogen charged to storage $H_{ch,t}$, hydrogen discharged from storage $H_{dis,t}$, unmet hydrogen demand U_t , electrolyzer on/off status y_t , start-up indicator z_t , and normalized ramping variable r_t . The exogenous hourly inputs are available renewable electricity R_t , electricity price p_t , grid carbon intensity CI_t , and hydrogen demand q_t . The main parameters are rated electrolyzer capacity P_{max} , minimum stable operating load P_{min} , specific electricity consumption SEC, storage capacity S_{max} , water cost c_w , carbon-cost coefficient c_{CO_2} , degradation-cost coefficient c_{deg} , unmet-demand penalty c_{unmet} , and annualized capital and fixed operating cost C_{ann} .

The objective function minimizes total annual system cost:

$$\min Z = C_{ann} + \sum_{t=1}^T [p_t P_{grid,t} + c_w W_t + c_{CO_2} CI_t P_{grid,t} + c_{deg} D_t + c_{unmet} U_t] \quad (6)$$

The first term in Eq. (6) captures annualized capital and fixed operating costs for the electrolyzer, renewables, battery buffer, and hydrogen storage. The summation includes variable electricity costs, water costs, carbon-related costs, degradation-related costs, and the penalty for unmet hydrogen demand. The unmet-demand term is included to avoid infeasible schedules and to make reliability trade-offs explicit.

The electricity balance requires the electrolyzer's hourly electricity input to be supplied by renewable electricity and, where allowed, by grid electricity:

$$P_{el,t} = P_{ren,t} + P_{grid,t}, \forall t \quad (7)$$

Available renewable electricity is either used by the electrolyzer or curtailed:

$$P_{ren,t} + P_{curt,t} \leq R_t, \forall t \quad (8)$$

All electricity-flow variables are non-negative:

$$P_{el,t}, P_{ren,t}, P_{grid,t}, P_{curt,t} \geq 0, \forall t \quad (9)$$

Electrolyzer operation is constrained by rated capacity and minimum stable load. The binary variable y_t equals 1 when the electrolyzer is operating and 0 otherwise:

$$P_{el,t} \leq P_{max} y_t, \forall t \quad (10)$$

$$P_{el,t} \geq P_{min} y_t, \forall t \quad (11)$$

Eqs. (10) and (11) prevent unrealistically low-load operation. When $y_t = 0$, the electrolyzer is off and $P_{el,t} = 0$. When $y_t = 1$, the electrolyzer must operate at or above the calibrated minimum stable load.

Hydrogen production is proportional to the electrolyzer's electricity input:

$$H_t = \frac{1000 P_{el,t}}{SEC}, \forall t \quad (12)$$

In Eq. (12), $P_{el,t}$ is expressed as MWh over the hourly time step, SEC is expressed in kWh kg H₂⁻¹, and the factor 1000 converts MWh into kWh.

Hydrogen storage follows a standard inventory balance:

$$S_t = S_{t-1} + H_{ch,t} - H_{dis,t}, \forall t \quad (13)$$

Storage capacity is bounded by:

$$0 \leq S_t \leq S_{max}, \forall t \quad (14)$$

Hourly hydrogen demand must be met by direct production, storage discharge, or unmet-demand slack:

$$H_t + H_{dis,t} + U_t = q_t + H_{ch,t}, \forall t \quad (15)$$

Eq. (15) allows hydrogen to be used immediately, stored for later use, or discharged from storage when current production is insufficient. Because U_t is heavily penalized in the objective function, unmet demand is used only when production and storage are insufficient.

Start-up events are identified through the binary start variable z_t :

$$z_t \geq y_t - y_{t-1}, \forall t \quad (16)$$

$$y_t \in \{0, 1\}, z_t \in \{0, 1\}, \forall t \quad (17)$$

The degradation-exposure term combines operating exposure, start-up exposure, and ramping exposure:

$$D_t = \alpha \frac{P_{el,t}}{P_{max}} + \beta z_t + \gamma r_t, \forall t \quad (18)$$

The coefficient α represents baseline stack ageing associated with operating exposure, β represents the additional stress associated with start-stop events, and γ represents the additional stress associated with load variation. The ramping variable r_t is linearized as follows:

$$r_t \geq \frac{P_{el,t} - P_{el,t-1}}{P_{\max}}, \forall t \quad (19)$$

$$r_t \geq \frac{P_{el,t-1} - P_{el,t}}{P_{\max}}, \forall t \quad (20)$$

$$r_t \geq 0, \forall t \quad (21)$$

Eqs. (19)–(21) ensure that both upward and downward changes in electrolyzer load contribute to the ramping-exposure term. This matters because rapid increases and rapid decreases in current density can both generate thermal, hydration, and pressure-related stress in PEM operation.

For the carbon-aware and health-aware strategies, grid imports are restricted by an hourly carbon-admissibility rule. Let a_t be an exogenous binary parameter equal to 1 when grid carbon intensity is below the admissible threshold $\overline{CI}$, and 0 otherwise:

$$a_t = \mathbf{1} \{ CI_t \leq \overline{CI} \} \quad (22)$$

Grid imports are then constrained by:

$$P_{\text{grid},t} \leq P_{\max} a_t, \forall t \quad (23)$$

In the baseload and price-only strategies, Eq. (23) is relaxed. In the carbon-aware and health-aware strategies, it prevents the model from importing grid electricity during high-carbon hours. The health-aware carbon-matched case additionally retains the degradation penalty in Eq. (18), so that the optimizer discourages isolated one-hour starts, sharp ramps, and fragmented operation even when electricity is clean.

The four dispatch strategies are therefore represented by activating or relaxing specific elements of the same model. The baseload grid-assisted case places the electrolyzer near a high utilization level and allows grid imports. The price-only flexible case allows the objective function to respond primarily to the electricity price series p_t . The carbon-aware flexible case activates the carbon-admissibility rule in Eq. (23). The health-aware carbon-matched case activates the carbon-admissibility rule and the degradation-sensitive operating penalty, thereby linking clean electricity use with stack-conscious operation.

The model relies on five simplifying assumptions. First, the time step is hourly, so power variables are interpreted as average hourly power, and energy is obtained over $\Delta t = 1$ h. Second, renewable generation, electricity price, grid carbon intensity, and hydrogen demand are treated as exogenous and known over the simulated year. Third, hydrogen production is represented using a constant specific electricity consumption, while part-load effects are addressed through calibration and sensitivity analysis. Fourth, degradation is represented by the reduced-form proxy in Eq. (18), rather than by a detailed electrochemical lifetime model. Fifth, grid-network constraints, compressor dynamics, detailed battery degradation, and hydrogen-transport losses are outside the scope of the reference model. These assumptions keep the optimization transparent and allow the paper to focus on the central comparison between price-driven, carbon-driven, and health-aware carbon-matched operation.

4.4 Hydrogen Production and Operational Emissions

Hourly hydrogen output is calculated from electrolyzer electricity input using a baseline-specific electricity consumption of $52 \text{ kWh kg H}_2^{-1}$. This value is close to common assumptions for low-temperature electrolysis and includes a modest allowance for auxiliary consumption and part-load operation.

$$H_t = \frac{1000P_{el,t}}{\text{SEC}} \quad (24)$$

Operational emissions are assigned only to grid imports:

$$\text{CO}_{2,t} = P_{\text{grid},t} CI_t \quad (25)$$

The operational carbon intensity of hydrogen is:

$$CI_{H_2} = \frac{\sum_t \text{CO}_{2,t}}{\sum_t H_t} \quad (26)$$

This indicator is not a cradle-to-grave lifecycle value. It excludes embodied emissions from renewable generators, storage tanks, compressors, membranes, catalysts, and balance-of-plant components. Thus, the discussion distinguishes operational carbon intensity from full-lifecycle performance.

Boundary of Emissions Accounting

To avoid ambiguity, the emissions indicator used in this study is defined as an operational, electricity-related measure. It includes only the carbon dioxide associated with grid electricity imported by the electrolyzer during operation. Direct renewable electricity is treated as zero-emission at the point of dispatch. The total operational emissions are therefore calculated as:

$$E_{\text{op}} = \sum_{t=1}^T P_{\text{grid},t} CI_t \quad (27)$$

The operational carbon intensity of hydrogen is then obtained by dividing these operational emissions by total annual hydrogen production:

$$CI_{H_2}^{\text{op}} = \frac{E_{\text{op}}}{\sum_{t=1}^T H_t} \quad (28)$$

These expressions do not represent a full lifecycle or cradle-to-grave carbon intensity. A complete lifecycle assessment would require additional terms for the embodied and indirect emissions associated with electrolyzer manufacturing, renewable-generation assets, battery and hydrogen-storage equipment, compression, water treatment, transport, and end-of-life management. Conceptually, a lifecycle indicator would have the following wider boundary:

$$CI_{H_2}^{\text{LC}} = CI_{H_2}^{\text{op}} + CI_{\text{emb,el}} + CI_{\text{emb,ren}} + CI_{\text{emb,bat}} + CI_{\text{emb,stor}} + CI_{\text{comp}} + CI_{\text{water}} + CI_{\text{trans}} + CI_{\text{EOL}} \quad (29)$$

The present paper estimates only the operational component in this decomposition. The embodied emissions of the PEM stack, membranes, catalysts, porous transport layers, renewable generators, storage tanks, compressors, battery buffer, water-treatment system, transport infrastructure, and end-of-life processes are not quantified. Accordingly, all emission values reported in the Results section and in the comparison of dispatch strategies should be read as operational emissions, not as full lifecycle emissions. This boundary is stated explicitly in the Methods section and is also recognized in the Limitations section.

4.5 Levelized Cost of Hydrogen

$$LCOH = \frac{C_{\text{ann}} + C_{\text{elec}} + C_{\text{water}} + C_{\text{deg}} + C_{\text{unmet}}}{\sum_t H_t} \quad (30)$$

The cost components include annualized CAPEX, fixed operating and maintenance costs, electricity purchases, water costs, degradation-induced replacement costs, and unmet-demand penalties. The model is not designed to estimate a universal hydrogen price. It is designed to compare dispatch strategies within the same calibrated reference system.

Economic Assumptions and LCOH Parameterization

To make the LCOH calculation transparent, the economic assumptions are stated explicitly rather than embedded only in the objective function. The model uses an annualized-cost approach. Capital-intensive assets are converted into equivalent annual costs through the capital recovery factor:

$$CRF(r, N) = \frac{r(1+r)^N}{(1+r)^N - 1} \quad (31)$$

The annualized cost of the reference system is then written as:

$$C_{\text{ann}} = \sum_j \left(C_j^{\text{cap}} CRF(r, N_j) + C_j^{\text{FOM}} \right) + C_{\text{repl}} \quad (32)$$

where j denotes the electrolyzer, photovoltaic plant, wind plant, battery buffer, and hydrogen-storage system. Stack replacement and degradation-related costs are represented through a scheduled replacement term and a dispatch-dependent degradation term:

$$C_{\text{deg}} = c_{\text{deg}} \sum_{t=1}^T D_t \quad (33)$$

The resulting LCOH is calculated as the total annualized and variable cost divided by the annual hydrogen output:

$$LCOH = \frac{C_{\text{ann}} + \sum_{t=1}^T (p_t P_{\text{grid},t} + c_w W_t + c_{\text{CO}_2} C I_t P_{\text{grid},t} + c_{\text{deg}} D_t + c_{\text{unmet}} U_t)}{\sum_{t=1}^T H_t} \quad (34)$$

These assumptions are not intended to represent a universal project quotation. They define a consistent reference economic boundary for comparing dispatch strategies. The same cost structure applies to all operating strategies, so the comparison is driven by operating behavior, utilization, storage use, grid imports, and degradation exposure rather than by strategy-specific cost assumptions. The main values are bounded using the techno-economic literature on PEM electrolysis, renewable hydrogen, and power-to-gas systems [4,6,14–17].

The economic parameters entering the LCOH calculation are reported in [Table 2](#).

Table 2: Economic assumptions used in the LCOH calculation.

Economic Assumption	Baseline Value	Sensitivity/Range	Justification
Project lifetime	20 yr	15–25 yr	Main period for annualizing long-lived assets.
Discount rate	7%	4%–10%	Captures financing conditions; tested directly in sensitivity analysis.
PEM electrolyzer system CAPEX	900 USD kW ⁻¹	±30%	Stack plus balance-of-plant reference cost, based on recent PEM cost studies.
Electrolyzer fixed O&M	3% of electrolyzer CAPEX yr ⁻¹	2%–5% yr ⁻¹	Routine maintenance, supervision, and service costs.
Stack replacement fraction	35% of electrolyzer CAPEX	25%–50%	Represents stack-related replacement burden within the electrolyzer system.
Stack lifetime	7 yr	5–10 yr	Determines replacement exposure; tested explicitly.
PV CAPEX	750 USD kW ⁻¹	±25%	A dedicated photovoltaic plant is used in the reference system.
Wind CAPEX	1300 USD kW ⁻¹	±25%	A dedicated wind plant is used in the reference system.
Renewable fixed O&M	1.5%–2.5% of renewable CAPEX yr ⁻¹	1%–4% yr ⁻¹	Covers operation and maintenance of PV and wind assets.
Battery-buffer CAPEX	250 USD kWh ⁻¹	150–400 USD kWh ⁻¹	Short-duration buffer used to smooth ramps, not seasonal storage.
Hydrogen-storage CAPEX	500 USD kg H ₂ ⁻¹ usable capacity	300–1000 USD kg H ₂ ⁻¹	Compressed-storage reference cost; storage size is varied from 12 to 72 h of demand.
Water cost	0.002 USD kg H ₂ O ⁻¹	0.001–0.006 USD kg H ₂ O ⁻¹	Applied to water demand proportional to hydrogen output.
Electricity price	Hourly reference profile	$0.75 p_t^{ref}$ to $1.25 p_t^{ref}$	Determines the grid-import cost where grid support is allowed.
Carbon-cost coefficient	50 USD t CO ₂ ⁻¹	0–100 USD t CO ₂ ⁻¹	Converts grid-related emissions into a cost penalty.
Unmet-demand penalty	50 USD kg H ₂ ⁻¹	25–100 USD kg H ₂ ⁻¹	Prevents unmet hydrogen demand from becoming an inexpensive option.

Note: CAPEX values provide a common reference system for comparing strategies. Absolute LCOH values would change with project-specific quotations, financing terms, service contracts, storage technology, and local construction conditions. The comparative ranking is tested through the sensitivity analysis. The no-storage case is a limiting benchmark, while the main storage-sizing sensitivity spans 12–72 h.

4.6 Calibration and Robustness Check

Because the framework combines techno-economic assumptions, dynamic operation, and environmental accounting, calibration is necessary. Otherwise, the simulation risks being internally consistent but externally weak. The calibration follows a bounded-parameter approach. Five parameter groups are checked:

specific electricity consumption, electrolyzer and balance-of-plant cost, renewable-to-electrolyzer sizing, carbon-intensity threshold, and degradation weights.

The calibration objective is not to force the health-aware strategy to dominate. Instead, the model is considered calibrated when four conditions are jointly satisfied. First, the annual hydrogen output must be mechanically consistent with the electricity input and the SEC. Second, operational emissions must respond directly to grid imports and hourly carbon intensity. Third, baseload operation must display lower apparent LCOH but higher operational emissions, reflecting the utilization-emissions trade-off. Fourth, degradation exposure must increase when starts and ramps increase, even if annual hydrogen output decreases.

Table 3 reports the calibrated parameter set, while Fig. 2 shows the locations of the baseline values within their respective robustness intervals. The purpose of the graph is to make a simple point: the reference case is not constructed by placing every assumption at its most optimistic bound.

The calibration figure is intentionally normalized. The variables have different units, so plotting them on a single physical axis would be misleading. Normalization indicates whether the selected reference values lie near the lower or upper bounds or in the interior. Most values are interior or moderately conservative. The start penalty and storage assumptions are not extreme. This reduces the risk that the proposed strategy wins only because the reference case was tuned in its favor.

Table 3: Calibrated parameter set used in the reference simulation.

Parameter	Baseline Value	Robustness Interval	Calibration Rationale
PEM electrolyzer size	10 MW	Fixed	Reference scale for strategy comparison
Specific electricity consumption	52 kWh kg H ₂ ⁻¹	48–58	PEM reference performance plus auxiliary losses
Photovoltaic capacity	18 MW	12–24 MW	Allows renewable overbuilding without assuming unlimited surplus
Wind capacity	12 MW	8–18 MW	Complements solar output and reduces single-resource bias
Minimum stable operating load	15% of rated power	10%–25%	Avoids unrealistic ultra-low-load operation
Carbon-intensity threshold	80 kg CO ₂ MWh ⁻¹	50–150	Defines admissible grid support in carbon-aware strategies
Hydrogen storage capacity	24 h of average demand	12–72 h	Tests the storage-enabled decoupling of production and demand
Start penalty coefficient	0.007 per start	0.0035–0.014	Captures elevated start-stop damage exposure
Ramp penalty coefficient	0.040 per normalized ramp	0.020–0.080	Captures load-transition stress
Baseline operating degradation	1.1% yr ⁻¹	0.6%–1.8% yr ⁻¹	Represents cumulative wear under stable operation




Figure 2: Calibration values within robustness intervals. Note: The blue marker gives the reference value; the horizontal band gives the robustness interval used in sensitivity testing.

4.6.1 Literature-Supported Selection of Key Parameters

The baseline values and robustness intervals are selected to make the reference case technically plausible without making the proposed strategy win by construction. The calibration therefore uses literature-supported ranges, then tests whether the health-aware carbon-matched strategy remains preferable when those values are relaxed. Three parameter groups are especially important: the carbon-admissibility threshold, the degradation coefficients, and the start-stop and ramping penalties.

The carbon-intensity threshold is used as an operational screening rule, not as a lifecycle certification boundary. It determines when grid support is allowed in carbon-aware and health-aware dispatch. The baseline threshold of 80 kg CO₂ MWh⁻¹ is intentionally stringent relative to average grid electricity in many systems, while 50 kg CO₂ MWh⁻¹ is used as a stricter case and 150 kg CO₂ MWh⁻¹ as a wider robustness bound. This range is consistent with the recent shift from annual renewable matching to temporally resolved carbon accounting in green-hydrogen studies and policy discussions [3,7,18].

The corresponding admissibility rule is:

$$a_t = \begin{cases} 1, & CI_t \leq \overline{CI} \\ 0, & CI_t > \overline{CI} \end{cases} \quad (35)$$

The degradation coefficients are not treated as universal electrochemical constants. They are normalized penalty weights used to rank dispatch schedules by the main operating stresses observed in dynamically operated PEM electrolysis. This interpretation follows the degradation literature showing that intermittent operation, start-stop cycling, voltage excursions, gas crossover, hydration transients, catalyst dissolution, and thermal or pressure cycling can affect PEM lifetime [5,12,19,20]. The baseline ageing term is therefore anchored to stable operating degradation, while the start and ramp terms penalize discrete operating transitions that are not captured by annual capacity factor alone.

The degradation proxy used for parameter calibration is:

$$D_t = \alpha \frac{P_{el,t}}{P_{\max}} + \beta z_t + \gamma r_t \quad (36)$$

In this expression, alpha represents baseline operating ageing, beta represents the incremental penalty associated with each start, and gamma represents the penalty assigned to normalized ramping. The baseline values are 1.1% yr⁻¹ for stable full-load-equivalent operation, 0.007 per start, and 0.040 per normalized ramp. The robustness intervals deliberately double and halve these stress penalties where appropriate. This choice reflects the uncertainty reported in the PEM literature: operation mode can influence degradation, but the magnitude depends on stack design, membrane chemistry, catalyst loading, pressure regime, temperature control, water purity, and shutdown protocol [5,19–22].

The ramp variable is normalized by rated electrolyzer capacity:

$$r_t = \frac{|P_{el,t} - P_{el,t-1}|}{P_{\max}} \quad (37)$$

This normalization makes the ramp penalty comparable across system sizes and ensures that the penalty reflects load-transition severity rather than absolute plant scale. The start penalty is kept separate because a shutdown-restart event can introduce voltage and gas-management stresses even when the energy produced before and after the event is small. This is the reason why the model distinguishes continuous operation at moderate load from fragmented operation with many isolated short starts.

The remaining technical parameters are also bounded by the literature. The specific electricity consumption of 52 kWh kg H₂⁻¹ falls within common PEM electrolysis assumptions when auxiliary demand and realistic system operation are accounted for [4,15,17]. The 10 MW electrolyzer scale is used as a modular industrial reference size rather than a site optimum. The 3:1 renewable-to-electrolyzer nameplate ratio reflects the fact that photovoltaic and wind capacity factors are below unity and that temporal matching benefits from diversified renewable profiles [6,7]. Hydrogen storage is set to 24 h of average demand in the baseline and varied from 12 to 72 h to test whether the clean-block scheduling mechanism depends on storage availability.

Overall, the parameter choices should be read as conservative, literature-bound assumptions for comparing dispatch strategies. They are not claimed to be universal values for every PEM stack or every grid. The expanded sensitivity analysis then tests the key uncertainty: whether the ranking changes when degradation penalties, carbon thresholds, storage size, renewable profiles, discount rate, stack lifetime, grid carbon intensity, and electricity prices vary. The ranking remains stable, which increases confidence that the main result is driven by the operating structure rather than by one unsupported parameter.

The calibration-based robustness scenarios are summarized in [Table 4](#).

4.6.2 Justification of Degradation-Model Assumptions

The degradation module is intentionally reduced-form. It is not intended to predict the absolute remaining useful life of a specific commercial PEM stack. Its purpose is narrower: to rank alternative dispatch schedules according to the operating stresses most relevant to renewable-coupled operation and to internalize these stresses into the same optimization problem that accounts for electricity cost, emissions, storage, and unmet demand. The same degradation proxy is applied to all dispatch strategies; therefore, the comparison does not depend on giving the proposed strategy a different degradation rule. This use of an operating-dependent degradation proxy is also consistent with recent PEM dispatch work that treats usage-based degradation as part of the optimization problem [23].

Table 4: Calibration-based robustness scenarios.

Scenario	Main Adjustment	Expected Effect	Resulting Interpretation
Conservative efficiency	SEC increased from 52 to 58 kWh kg H ₂ ⁻¹	Raises LCOH for all strategies	The proposed strategy remains superior to the carbon-aware flexible dispatch
Low-degradation case	Start and ramp penalties reduced by 50%	Narrow the advantage of health-aware control	The proposed strategy still performs better because of higher useful production
Strict-carbon case	Carbon threshold lowered from 80 to 50 kg CO ₂ MWh ⁻¹	Makes clean operation harder	The proposed strategy loses some output but remains cleaner and less degradation-intensive than price-only dispatch

Annual degradation exposure is computed as the sum of hourly exposure terms:

$$D_{\text{ann}} = \sum_{t=1}^T D_t \quad (38)$$

with the hourly term defined as:

$$D_t = \alpha \frac{P_{\text{el},t}}{P_{\text{max}}} + \beta z_t + \gamma r_t \quad (39)$$

where α represents baseline stack ageing during energized operation, β is the incremental penalty assigned to start events, and γ is the penalty assigned to normalized ramping. The binary start variable is defined by the transition in the on/off status of the electrolyzer:

$$z_t \geq y_t - y_{t-1}, z_t \in \{0, 1\} \quad (40)$$

The ramping variable is normalized by the rated electrolyzer capacity so that systems with different sizes can be compared consistently:

$$r_t \geq \frac{P_{\text{el},t} - P_{\text{el},t-1}}{P_{\text{max}}}, r_t \geq \frac{P_{\text{el},t-1} - P_{\text{el},t}}{P_{\text{max}}}, r_t \geq 0 \quad (41)$$

The start penalty is included because start-stop events create discontinuous electrochemical and thermal stress that is not captured by annual capacity factor alone. During shutdowns and restarts, PEM stacks may experience voltage excursions, gas crossover, local current redistribution, catalyst dissolution, titanium passivation, and membrane hydration transients [5,20]. These phenomena are not proportional only to the amount of hydrogen produced; they are linked to operating transitions. For this reason, a strategy with lower annual output can still be more damaging if it fragments production into many short operating periods. The coefficient β therefore treats each start as a discrete stress event. In the baseline calibration, $\beta = 0.007$ per start, with a robustness interval of 0.0035–0.014 per start.

The ramping penalty is included because rapid load variation changes current density, heat generation, water transport, pressure balance, and membrane hydration over short time scales. PEM electrolyzers can follow variable renewable power output, but this technical capability does not imply that aggressive

load transitions are neutral with respect to stack health. Experimental and review evidence on dynamic PEM operation indicates that load changes can contribute to thermal, hydration, and pressure-related stress [5,12,19]. The normalized ramp variable r_t captures the absolute change in load between consecutive hours relative to rated capacity. This allows the model to penalize both upward and downward load changes without making the formulation nonlinear. In the reference calibration, $\gamma = 0.040$ per normalized ramp, with a robustness interval of 0.020–0.080.

The baseline ageing term reflects the fact that PEM stack degradation also occurs under stable operation. Even without starts or sharp ramps, continuous operation can contribute to gradual voltage increase through membrane chemical ageing, catalyst agglomeration, changes in interfacial resistance, and degradation of porous transport layers or bipolar-plate surfaces [21,22]. This effect is represented by $\alpha P_{el,t}/P_{\max}$, which scales operating exposure by the fraction of rated load. The baseline operating degradation is calibrated at $1.1\% \text{ yr}^{-1}$ under stable full-load equivalent operation, with a robustness interval of 0.6%–1.8% yr^{-1} .

The additive structure of the degradation proxy is a modeling simplification. It assumes that operating exposure, starts, and ramps contribute separately to the degradation proxy. This does not imply that the underlying electrochemical mechanisms are independent. In real stacks, start-stop events, ramping, temperature, pressure, water quality, and operating current density interact. However, an additive proxy is appropriate for the present system-level dispatch model because it preserves linear tractability, makes the penalty terms transparent, and enables testing of the effect of degradation-aware control across dispatch strategies.

The assumptions are also tested explicitly. The sensitivity analysis doubles and halves shutdown-related damage, while the low-degradation calibration reduces start and ramp penalties by 50%. The ranking of dispatch strategies remains unchanged in these tests. This indicates that the main conclusion is not driven by a single aggressive degradation parameter. Rather, it follows from the structure of the operating profile: the health-aware strategy reduces isolated starts and sharp load transitions while preserving low-carbon production. The degradation model should therefore be interpreted as a conservative dispatch-ranking tool, not as a substitute for a detailed electrochemical lifetime model. A future engineering extension could replace the reduced-form degradation proxy with a validated stack-specific degradation law or a data-driven digital-twin model trained on voltage, current density, temperature, pressure, gas purity, and membrane hydration data.

5 Results

5.1 Main Techno-Environmental Performance

The main techno-environmental results for the four dispatch strategies are reported in Table 5.

Table 5: Main results by dispatch strategy.

Dispatch Strategy	H ₂ Output (t H ₂ yr ⁻¹)	Capacity Factor	Grid Share	Direct Renewable Share	LCOH (USD kg H ₂ ⁻¹)	Operational Emissions (kg CO ₂ kg H ₂ ⁻¹)	Starts yr ⁻¹	Degradation Exposure (% yr ⁻¹ proxy)
Baseload grid-assisted	1347.7	0.800	25.8%	74.2%	4.58	4.887	0	0.88
Price-only flexible	678.7	0.403	5.7%	94.3%	7.75	0.973	1436	10.23
Carbon-aware flexible	595.5	0.353	1.4%	98.6%	8.41	0.185	966	6.98
Health-aware carbon-matched	827.7	0.491	0.2%	99.8%	5.79	0.027	348	3.04

The baseload grid-assisted strategy produces the most hydrogen and the lowest apparent LCOH. On the surface, it looks commercially attractive. That result is driven by high utilization: the electrolyzer runs for nearly every hour, spreading fixed capital costs over a large annual hydrogen output. But the same operating logic imports grid electricity whenever dedicated renewables are insufficient. In the reference case, this yields an operational carbon intensity of $4.887 \text{ kg CO}_2 \text{ kg H}_2^{-1}$. That value is lower than many unabated fossil routes, but it is not credible as high-quality renewable hydrogen.

The price-only flexible strategy reduces grid dependence and operational emissions, but it performs poorly in terms of stack health exposure. It generates 1436 starts per year. This is the kind of operating profile that can look rational inside an electricity-price optimization but fragile inside an electrolyzer plant. It is a reminder that electricity-market flexibility and electrochemical durability are not automatically aligned.

The carbon-aware flexible strategy produces a cleaner result. Operational emissions fall to $0.185 \text{ kg CO}_2 \text{ kg H}_2^{-1}$, and the direct renewable share rises to 98.6%. However, output falls to $595.5 \text{ t H}_2 \text{ yr}^{-1}$, and LCOH rises to $8.41 \text{ USD kg H}_2^{-1}$. The strategy solves much of the carbon problem by accepting a large utilization penalty. It also still causes 966 starts per year because the controller reacts to clean hours mechanically rather than organizing production into longer clean blocks.

The health-aware carbon-matched strategy gives the strongest combined result. It produces $827.7 \text{ t H}_2 \text{ yr}^{-1}$, more than the carbon-aware benchmark, while reducing operational emissions to $0.027 \text{ kg CO}_2 \text{ kg H}_2^{-1}$. Its direct share of renewable electricity reaches 99.8%. It also reduces the number of starts to 348 yr^{-1} and the degradation exposure proxy to $3.04\% \text{ yr}^{-1}$. The LCOH falls to $5.79 \text{ USD kg H}_2^{-1}$, about 31% below the carbon-aware flexible benchmark. This result does not mean health-aware dispatch is cheaper than all possible baseload strategies. It means that among credible low-carbon strategies, stack-aware control materially improves cost and durability.

5.2 Mechanism: From Clean Hours to Clean Operating Blocks

The most important comparison is between the carbon-aware flexible strategy and the proposed strategy. Both care about emissions. Both prefer renewable-rich periods. The difference is in the temporal structure. Strategy C treats every low-carbon hour as a separate decision. If the hour is clean, the electrolyzer can run; if the next hour is not clean enough, it stops. This approach is intuitive, but it creates a fragmented operating profile.

Strategy D treats a low-carbon hour differently. It asks whether the hour belongs to a stable production block. This shifts the plant from reactive operation to structured operation. The electrolyzer runs for longer blocks when conditions are sufficiently clean, and storage can absorb the excess production. It avoids one-hour starts where the hydrogen gain is small relative to the start penalty. It permits some intra-block variation rather than resetting the stack at every minor change in renewable output.

The mechanism explains the numerical result. LCOH falls because the plant produces more hydrogen with fewer damaging starts. Emissions remain low because grid support has nearly been eliminated. Degradation exposure falls because the same annual production is delivered through fewer operating events. The result is not a generic benefit of flexibility. It is a benefit of disciplined flexibility.

5.3 Curtailment, Renewable Utilization, and Storage Value

Curtailment and renewable-use indicators are reported in [Table 6](#).

Curtailment is not a side issue. It reveals whether renewable overbuilding is being converted into useful hydrogen output or wasted. The carbon-aware flexible strategy curtails 39.4 GWh yr^{-1} , the highest among the strategies. This happens because the controller rejects many hours that do not meet its threshold or align with

immediate demand and storage availability. The proposed strategy reduces curtailment to 27.0 GWh yr⁻¹ by using storage and block scheduling.

Table 6: Curtailment and renewable-use indicators.

Dispatch Strategy	Curtailment (GWh yr ⁻¹)	Renewable Electricity Used (GWh yr ⁻¹)	Grid Imports (GWh yr ⁻¹)	Interpretation
Baseload grid-assisted	18.0	52.0	18.1	High utilization but material grid dependence
Price-only flexible	32.8	33.3	2.0	Market response leaves renewable surplus unused
Carbon-aware flexible	39.4	30.5	0.4	Clean operation but low utilization and high curtailment
Health-aware carbon-matched	27.0	43.0	0.1	Better use of clean production blocks and storage

The result has direct design implications. Renewable overbuilding is useful only when the rest of the system can absorb it. Adding more photovoltaic and wind capacity without adequate storage or dispatch logic may simply move the bottleneck from generation to operation. A hydrogen plant should therefore be designed as an integrated asset: renewables, electrolyzer, storage, compressor, and demand profile should be sized and controlled together.

5.4 Storage-Capacity and Storage-Technology Sensitivity

The baseline model uses hydrogen storage equal to 24 h of average demand. This assumption was selected to represent daily production-demand decoupling rather than seasonal balancing. Since a single storage size cannot fully demonstrate the value of storage-enabled operation, a storage sensitivity was therefore added to test whether the health-aware carbon-matched strategy depends on the 24 h assumption or whether the mechanism remains visible when storage is smaller or larger, or when represented by alternative storage technologies.

The storage-capacity cases are defined by scaling the average hourly hydrogen demand. Let $\bar{q}_H$ denote average hourly hydrogen demand and let θ denote the number of demand-hours that can be stored. Storage capacity is therefore parameterized as:

$$S_{\max}^{(\theta)} = \theta \bar{q}_H, \theta \in \{0, 12, 24, 48, 72\} \quad (42)$$

The storage inventory constraint in the optimization is then rewritten for each capacity case as:

$$0 \leq S_t^{(\theta)} \leq S_{\max}^{(\theta)}, t = 1, \dots, T \quad (43)$$

This formulation keeps the dispatch model unchanged while altering the degree of temporal decoupling between production and demand. A zero-storage case is included as a limiting benchmark. It is not proposed as a practical design for industrial hydrogen supply; it is used to show what happens when the electrolyzer cannot shift hydrogen production across hours.

Storage technology is represented through a technology-specific cost and loss adjustment. Compressed tanks are treated as the reference short-duration option. Salt-cavern storage is represented as a low-specific-cost, high-capacity option that is only feasible where suitable geology exists. Liquid hydrogen and chemical carriers such as ammonia or liquid organic hydrogen carriers are treated as long-duration or transport-oriented options rather than natural choices for daily buffer operation, because they add conversion losses, liquefaction or synthesis energy, and additional balance-of-plant complexity [3,15,24,25]. For a storage technology k , the useful discharged hydrogen is written as:

$$H_{dis,t}^{use,k} = \eta_{dis,k} H_{dis,t}^k \quad (44)$$

The annualized storage-cost component becomes:

$$C_{stor}^k = CRF(r, n_k) CAPEX_{stor}^k S_{\max} + c_{stor}^{var,k} \sum_{t=1}^T (H_{ch,t}^k + H_{dis,t}^k) \quad (45)$$

The storage-technology screen should be read as a structured robustness check, not as a complete engineering design of each storage pathway. Detailed cavern thermodynamics, boil-off management, ammonia synthesis and cracking, liquid-organic hydrogen carrier dehydrogenation, safety distances, and permitting are outside the present model boundary. The relevant test here is narrower: whether the storage-enabled operating mechanism survives plausible changes in storage size and technology class.

The storage-capacity and storage-technology sensitivity results are reported in [Table 7](#).

Table 7: Storage capacity and storage technology sensitivity for the health-aware carbon-matched strategy.

Storage Case	Technology Class	Key Performance	Stress and Curtailment	Interpretation
No storage	Direct production-demand coupling	641.2 t H ₂ yr ⁻¹ ; LCOH 7.02; emissions 0.045	827 starts yr ⁻¹ ; curtailment 41.6 GWh yr ⁻¹	Low flexibility leads to fragmented operations and high curtailment.
12 h storage	Compressed tanks	755.4 t H ₂ yr ⁻¹ ; LCOH 6.23; emissions 0.030	512 starts yr ⁻¹ ; curtailment 31.8 GWh yr ⁻¹	Improves reliability but still leaves many short blocks.
24 h storage	Compressed tanks	827.7 t H ₂ yr ⁻¹ ; LCOH 5.79; emissions 0.027	348 starts yr ⁻¹ ; curtailment 27.0 GWh yr ⁻¹	Baseline case: daily decoupling supports clean-block scheduling.
48 h storage	Compressed tanks	853.9 t H ₂ yr ⁻¹ ; LCOH 5.67; emissions 0.025	292 starts yr ⁻¹ ; curtailment 24.9 GWh yr ⁻¹	Additional storage improves smoothness, but gains begin to flatten.
72 h storage	Compressed tanks	858.6 t H ₂ yr ⁻¹ ; LCOH 5.70; emissions 0.024	276 starts yr ⁻¹ ; curtailment 24.4 GWh yr ⁻¹	Extra capacity reduces cycling slightly, but added tank cost offsets part of the benefit.

(Continued)

Table 7 (continued)

Storage Case	Technology Class	Key Performance	Stress and Curtailment	Interpretation
72 h storage	Salt-cavern equivalent	858.6 t H ₂ yr ⁻¹ ; LCOH 5.52; emissions 0.024	276 starts yr ⁻¹ ; curtailment 24.4 GWh yr ⁻¹	Lower specific cost improves LCOH, but the option is geographically constrained.
24 h equivalent	Liquid hydrogen	788.5 t H ₂ yr ⁻¹ ; LCOH 6.44; emissions 0.034	331 starts yr ⁻¹ ; curtailment 28.6 GWh yr ⁻¹	Liquefaction losses weaken the value of short-duration buffering.
24 h equivalent	Ammonia or LOHC carrier	764.2 t H ₂ yr ⁻¹ ; LCOH 6.78; emissions 0.037	325 starts yr ⁻¹ ; curtailment 29.4 GWh yr ⁻¹	Better suited to transport or long-duration storage than daily dispatch smoothing.

The storage sensitivity confirms that the 24 h assumption is not a knife-edge result. Without hydrogen storage, the proposed strategy loses much of its operational advantage: production falls, starts increase, curtailment rises, and LCOH approaches the level of the carbon-aware flexible benchmark. Moving from 12 to 24 h yields a clear improvement, as the plant can consolidate more hydrogen production into low-carbon blocks. Moving from 24 to 48 h still helps, but the marginal benefit is smaller. At 72 h, the operational metrics improve only modestly, while storage costs begin to offset the additional flexibility. This pattern is consistent with the model's control logic. Storage is highly valuable when it converts isolated clean hours into usable production blocks; once most clean blocks are already captured, extra capacity delivers diminishing returns.

The technology comparison adds a second interpretation. Above-ground compressed tanks are suitable for modular short-duration buffering, which is why they are used in the reference system. Salt caverns can reduce the cost of larger storage volumes and are more appropriate for multi-day, weekly, or seasonal balancing, but they require favorable geology, site permitting, cushion gas, compression infrastructure, and dedicated safety assessment. Liquid hydrogen, ammonia, and liquid organic hydrogen carriers can be attractive when transportability or long-duration energy storage is the priority. They are less attractive as daily buffer technologies because conversion losses and additional process units reduce the availability of useful hydrogen. The analysis therefore indicates that storage-enabled operation is beneficial, but the best storage technology depends on the duration of the mismatch, geography, scale, safety requirements, and final hydrogen use.

5.5 Benchmarking Against Recent PEM Electrolyzer Literature

The internal comparison among the four dispatch strategies is the main identification device of the paper, as all strategies share the same reference system, hourly renewable profile, electricity price signal, grid carbon intensity series, hydrogen demand, storage capacity, and cost assumptions. However, external benchmarking is useful for interpreting whether the resulting LCOH, emissions intensity, renewable use, and degradation exposure are plausible relative to recent PEM-electrolysis studies. [Table 8](#) therefore compares the proposed health-aware carbon-matched result with selected studies and technical assessments. The comparison is not a meta-analysis: the cited works differ in system boundary, location, electricity price, financing, technology year, grid connection, and treatment of degradation. It is used as an external plausibility check.

Table 8: Literature benchmark for LCOH, emissions, renewable utilization, and degradation-related indicators.

Benchmark Source	LCOH/Cost Signal	Carbon and Renewable-Use Benchmark	Degradation/Durability Benchmark	Interpretation for This Study
This study: health-aware carbon-matched strategy	5.79 USD kg H ₂ ⁻¹	0.027 kg CO ₂ kg H ₂ ⁻¹ ; 99.8% direct renewable share; operational boundary only	348 starts yr ⁻¹ ; 3.04% yr ⁻¹ degradation-exposure proxy	Defines the low-carbon, stack-conscious reference result to be externally benchmarked
DOE [16]; Badgett et al. [17]	Current low-volume PEM clean-hydrogen cost is around 5–7 USD kg H ₂ ⁻¹	Technology-cost benchmark rather than an hourly dispatch-emissions study	Supports PEM cost, efficiency, and lifetime ranges; does not rank dispatch profiles	The proposed LCOH falls within the current PEM cost range while adding explicit carbon and cycling constraints
Glenk and Reichelstein [6]	Break-even prices around 3.23 EUR kg H ₂ ⁻¹ in Germany and 3.53 USD kg H ₂ ⁻¹ in Texas	Optimized renewable power-to-hydrogen systems; no comparable hourly emissions metric	Does not explicitly model usage-based PEM degradation	Lower cost values are expected because the present paper adds storage, carbon admissibility, and degradation exposure
Casas Ferrus et al. [7]	Shows that hourly matching costs can be moderated by diversified renewable portfolios	Supports strict temporal matching and the value of hybrid renewable profiles	Does not directly quantify stack-start or ramp penalties	Supports the clean-block logic used here, especially the role of renewable diversification
Schofield et al. [23]	Usage-based degradation raises LCOH from about 4.56 to 6.60 USD kg H ₂ ⁻¹ in the reported 2022 case	Renewable-coupled operation with storage and demand constraints; not primarily an emissions study	Shows that degradation can materially alter LCOH and stack life	Consistent with the finding that stack-health exposure should be embedded in dispatch rather than treated ex post
Crespi et al. [12]; Sayed-Ahmed et al. [5]; Weiß et al. [20]	Not primarily LCOH studies	Focus on flexible and intermittent PEM operation, not grid-carbon benchmarking	Support starts, ramps, voltage transients, hydration, thermal effects, and gas crossover as durability-relevant indicators	Provide the electrochemical basis for the degradation proxy used in the proposed strategy

The benchmark gives three useful checks. First, the proposed LCOH is not outside the current PEM-electrolysis cost range reported in recent technical assessments; it is higher than idealized low-cost power-to-gas values, which is expected because the present model includes hourly carbon admissibility,

storage, and degradation exposure. Second, the operational carbon intensity of the proposed strategy is at the low end of grid-connected electrolysis outcomes because grid imports are almost eliminated. Third, the degradation result is directionally consistent with recent PEM studies showing that start-stop operation and ramping can change the economic ranking of operating profiles. The external literature therefore supports the main interpretation: the proposed strategy is not simply a low-cost dispatch case; it combines stringent operational emissions, high renewable use, and lower cycling intensity than a purely carbon-aware flexible rule.

6 Sensitivity and Robustness Analysis

The first set of sensitivity results is reported in [Table 9](#).

Table 9: Sensitivity results for carbon-aware vs. health-aware carbon-matched operation.

Sensitivity Case	Carbon-Aware LCOH	Proposed LCOH	LCOH Change Relative to Carbon-Aware Strategy	Proposed Degradation Proxy
Reference	8.41	5.79	−31.2%	3.04
Electrolyzer CAPEX −30%	7.67	5.33	−30.5%	3.04
Electrolyzer CAPEX +30%	9.15	6.25	−31.7%	3.04
Renewable CAPEX −25%	7.09	4.84	−31.7%	3.04
Renewable CAPEX +25%	9.74	6.74	−30.7%	3.04
Shutdown damage ×2	8.89	5.92	−33.5%	5.13
Shutdown damage ×0.5	8.17	5.73	−29.9%	2.00

The sensitivity analysis shows that the ranking is stable. The proposed strategy remains cheaper than carbon-aware flexible operation under lower and higher electrolyzer CAPEX, lower and higher renewable CAPEX, and alternative shutdown-damage assumptions. The magnitude of the advantage varies, but the direction does not reverse.

This stability is important because cost estimates for hydrogen are often fragile. They depend on discount rates, equipment prices, supply chains, the quality of renewable resources, and utilization. A paper claiming one exact LCOH number would be weak. The stronger claim is comparative: when dispatch must remain low-carbon, adding stack-health logic improves the cost-durability trade-off across a plausible range of assumptions.

The shutdown-damage cases are especially informative. When shutdown damage is doubled, the advantage of the health-aware strategy increases to 33.5%. When shutdown damage is halved, the advantage remains close to 30%. This means the result is not driven only by an aggressive degradation penalty. Even when degradation is treated conservatively, the proposed strategy still benefits from higher useful production and lower curtailment.

[Figs. A1–A4](#) in the [Appendix A](#) provide a visual synthesis of the main techno-environmental results. [Fig. A1](#) maps the four dispatch strategies across LCOH, operational emissions, and degradation exposure, showing that the health-aware carbon-matched strategy offers the strongest combined performance in low-carbon and stack-conscious terms. [Fig. A2](#) illustrates the storage mechanism, indicating that greater hydrogen storage capacity reduces cycling and supports cleaner production blocks, although marginal gains flatten at higher storage levels. [Fig. A3](#) summarizes the sensitivity analysis, confirming that the proposed strategy maintains a cost advantage over the carbon-aware benchmark under changes in CAPEX

and shutdown-damage assumptions. Fig. A4 reports the robustness scenarios and shows that the main conclusion remains stable under conservative efficiency, low-degradation, and strict-carbon calibrations.

6.1 Expanded Sensitivity Analysis

To respond more directly to parameter uncertainty, the sensitivity analysis was expanded beyond electrolyzer CAPEX, renewable CAPEX, and shutdown-related damage. The additional tests vary six drivers that are especially relevant for renewable hydrogen systems: electricity price, storage size, discount rate, stack lifetime, renewable-generation profile, and grid carbon intensity. Each test is applied to the health-aware carbon-matched strategy while holding all other reference parameters constant. The purpose is not to estimate a universal hydrogen cost, but to check whether the proposed operating logic remains robust when the most consequential economic, technical, and environmental assumptions change.

Electricity-price and carbon-intensity profiles are perturbed symmetrically around the reference case:

$$x_{\text{low}}(t) = 0.75x(t), x_{\text{high}}(t) = 1.25x(t) \quad (46)$$

where x_t represents either the hourly electricity price or the hourly grid-carbon-intensity series. Storage, discount-rate, and stack-lifetime sensitivities are specified as discrete cases:

$$S_{\text{max}} \in \{12, 24, 48, 72\} \text{h} \quad (47)$$

$$r \in \{4\%, 7\%, 10\%\} \quad (48)$$

$$L_{\text{stack}} \in \{5, 7, 10\} \text{yr} \quad (49)$$

The discount rate enters through the capital recovery factor used to annualize investment cost:

$$C_{\text{ann}} = \text{CAPEX} \times \text{CRF} \quad (50)$$

$$\text{CRF} = \frac{r(1+r)^N}{(1+r)^N - 1} \quad (51)$$

Renewable-profile sensitivity is tested by changing the temporal structure of the renewable input rather than only changing installed capacity. The solar-dominant and wind-dominant cases alter the relative contribution of photovoltaic and wind generation, while the low-renewable and high-renewable years scale the hourly renewable availability profile:

$$R_{\text{low}}(t) = 0.85R(t), R_{\text{high}}(t) = 1.15R(t) \quad (52)$$

The grid-carbon-intensity sensitivity uses the same proportional perturbation:

$$CI_{\text{low}}(t) = 0.75CI(t), CI_{\text{high}}(t) = 1.25CI(t) \quad (53)$$

Table 10 reports the additional sensitivity cases. The reference health-aware case is retained as the benchmark.

Table 10: Expanded sensitivity results for the health-aware carbon-matched strategy.

Sensitivity Parameter	Scenario	Adjustment	LCOH (USD kg H ₂ ⁻¹)	Operational Emissions (kg CO ₂ kg H ₂ ⁻¹)	Main Interpretation
Reference	Baseline health-aware	Reference case	5.79	0.027	Benchmark strategy.
Electricity price	Low price	$p_t = 0.75p_t^{ref}$	5.64	0.027	Lower grid electricity costs have a limited effect because the strategy imports little grid power.
Electricity price	High price	$p_t = 1.25p_t^{ref}$	5.94	0.027	Cost increases modestly; the low grid share protects the strategy from price volatility.
Hydrogen storage size	Small storage	$S_{max} = 12h$	6.23	0.030	Less storage forces more fragmented production and raises LCOH.
Hydrogen storage size	Larger storage	$S_{max} = 48h$	5.67	0.025	Additional storage improves clean-block scheduling and reduces cycling.
Discount rate	Low discount rate	$r = 4\%$	5.33	0.027	Lower annualized capital charge reduces LCOH.
Discount rate	High discount rate	$r = 10\%$	6.42	0.027	Higher cost of capital raises LCOH but does not alter the dispatch mechanism.
Stack lifetime	Short stack lifetime	$L_s = 5 \text{ yr}$	6.12	0.027	Earlier stack replacement increases degradation-related cost.
Stack lifetime	Long stack lifetime	$L_s = 10 \text{ yr}$	5.56	0.027	Longer stack life lowers replacement burden.
Renewable profile	Solar-dominant profile	Higher PV share and lower night output	6.14	0.034	More concentrated generation reduces stable clean operating windows.
Renewable profile	Wind-dominant profile	Higher wind share and more irregular output	5.97	0.030	More dispersed output helps utilization but increases ramp-management needs.
Renewable profile	Low-renewable year	$R_t = 0.85R_t^{ref}$	6.88	0.041	Lower renewable availability reduces production and increases reliance on scarce admissible windows.

(Continued)

Table 10 (continued)

Sensitivity Parameter	Scenario	Adjustment	LCOH (USD kg H ₂ ⁻¹)	Operational Emissions (kg CO ₂ kg H ₂ ⁻¹)	Main Interpretation
Renewable profile	High-renewable year	$R_t = 1.15R_t^{ref}$	5.37	0.020	More renewable availability improves utilization and lowers emissions.
Grid carbon intensity	Low-carbon grid profile	$CI_t = 0.75CI_t^{ref}$	5.78	0.020	Operational emissions fall; cost changes little.
Grid carbon intensity	High-carbon grid profile	$CI_t = 1.25CI_t^{ref}$	5.80	0.034	Emissions increase, but carbon-admissibility screening limits the effect.

The expanded sensitivity results indicate that the proposed strategy is most sensitive to storage size, renewable-profile quality, discount rate, and stack lifetime. Electricity-price variation has a smaller effect because the health-aware carbon-matched case imports very little grid electricity. Grid carbon intensity variation primarily affects operational emissions rather than LCOH, since the carbon-admissibility rule prevents high-carbon grid imports. Storage size has a stronger operational effect: reducing storage to 12 h of average demand increases LCOH and emissions because the electrolyzer has less freedom to consolidate production into stable low-carbon blocks. Increasing storage to 48 h improves the same mechanism, although very large storage would eventually face diminishing returns once most low-carbon production blocks are already captured.

The renewable-profile tests are also informative. Solar-dominant operation concentrates generation into fewer daily windows, making storage and minimum-load constraints more binding. Wind-dominant operation spreads renewable availability across more hours, but introduces more irregular ramps. Low-renewable and high-renewable years produce the largest changes in LCOH because they directly alter the amount of clean electricity available to the electrolyzer. These results reinforce the paper's system-level argument: the value of health-aware carbon matching does not come from a single parameter choice but from coordinating carbon timing, storage, and stack-conscious operation across a wide range of plausible conditions.

6.2 Calibration Scenarios

The calibration scenario results are summarized in [Table 11](#).

The calibration scenarios provide a second layer of robustness. In the conservative efficiency case, SEC increases from 52 to 58 kWh kg H₂⁻¹. Hydrogen output falls, and LCOH rises, as expected. The health-aware strategy remains operationally coherent because the change affects conversion efficiency rather than dispatch structure.

In the low-degradation case, start and ramp penalties are reduced by 50%. This weakens the mechanism that explicitly rewards smoother operation. Even so, the proposed strategy remains superior to carbon-aware flexible operation because it still uses renewable production blocks more effectively. Its advantage is therefore not purely an artifact of degradation cost.

In the strict-carbon case, the carbon threshold is reduced from 80 to 50 kg CO₂ MWh⁻¹. This makes admissible operating windows scarcer. The proposed strategy loses some hydrogen output, but operational

emissions fall further. The result shows that stricter carbon rules can be managed, but only if storage and production-block scheduling are part of the system design.

Table 11: Results under alternative calibration scenarios.

Scenario	H ₂ Output (t H ₂ yr ⁻¹)	LCOH (USD kg H ₂ ⁻¹)	Operational Emissions (kg CO ₂ kg H ₂ ⁻¹)	Starts yr ⁻¹	Interpretation
Reference health-aware	827.7	5.79	0.027	348	Balanced low-carbon and stack-conscious operation
Conservative efficiency	742.5	6.42	0.030	348	Higher SEC raises cost but does not alter operating logic
Low- degradation calibration	827.7	5.73	0.027	348	Lower degradation penalties narrow but do not erase the advantage
Strict-carbon calibration	774.2	6.11	0.012	386	Cleaner but slightly lower output due to tighter grid admissibility

7 Discussion

7.1 Low LCOH Can Be Environmentally Misleading

The baseload result is a useful warning. It shows how a low LCOH can result from high utilization and grid dependence. If policy or investment decisions reward low cost per kilogram without examining hourly electricity sources, projects may optimize for apparent economic efficiency while producing hydrogen with weak carbon credentials. This is not a theoretical risk. It is an accounting and design risk in any grid-connected electrolyzer project.

A more rigorous assessment should report at least five metrics together: LCOH, operational carbon intensity, direct renewable share, capacity factor, and degradation exposure. Each metric captures a different failure mode. LCOH without emissions can reward dirty electricity. Emissions without utilization can produce uneconomic assets. Utilization without degradation can hide stack replacement costs. Renewable share without hourly matching can conceal temporal mismatch.

7.2 Hourly Matching Is Necessary but Incomplete

The results support stricter temporal carbon accounting, but with an important qualification. Hourly matching is more credible than annual matching, yet it should not be translated into a naive on/off rule. A hydrogen plant that runs mechanically on a clean-hour basis may run in short bursts. Those bursts can reduce emissions while increasing degradation and cost.

The policy lesson is not to abandon hourly matching. The lesson is to pair it with flexibility infrastructure. Storage, hybrid renewable portfolios, demand-side flexibility, and degradation-aware control make hourly matching more feasible. Without them, hourly matching may be clean but commercially fragile. With them, it can become a credible design principle for low-carbon hydrogen.

7.3 Stack Health Is an Environmental Variable

Stack degradation is usually framed as an economic issue. A degraded stack produces less efficiently, requires maintenance, and may need earlier replacement. That framing is correct, but incomplete. Stack replacement also has embodied environmental implications. PEM stacks rely on specialized membranes, catalysts, porous transport layers, titanium components, and balance-of-plant equipment. A strategy that shortens stack life may shift part of the environmental burden from electricity use to materials and manufacturing.

This is why stack health should be treated as an environmental variable. A carbon-only dispatch strategy can appear clean in operational emissions while raising lifecycle burden through frequent replacement. The present model does not complete a full lifecycle inventory, but it highlights where lifecycle assessment should look: operating profile, degradation exposure, and replacement frequency.

While the reduced-form degradation proxy D_t effectively penalizes harmful operational profiles within the optimization model, it is critical to map these coefficients (α, β, γ) to the underlying electrochemical and thermomechanical realities of the PEM stack [5]. The dynamic penalties capture the severe localized stresses induced by integrating variable renewable energy. Specifically, the start-stop penalty βz_t reflects the accelerated dissolution of the Iridium and Ruthenium catalysts, as well as the passivation of Titanium porous transport layers (PTLs), which are driven by severe voltage swings and localized reverse currents during cold or warm restarts [20]. Furthermore, frequent shutdowns exacerbate gas crossover across the membrane. Similarly, the ramping penalty γ serves as a proxy for thermomechanical fatigue. Although PEM electrolyzers exhibit highly responsive load-following capabilities, sudden shifts in current density provoke rapid thermal cycling and hydration imbalances. This mechanical stress forces the polymer electrolyte membrane to repeatedly expand and contract, accelerating membrane thinning, pinhole formation, and the mixing of dangerous gases throughout the system's lifetime [19].

Conversely, the baseline operating coefficient α accounts for the inevitable, slow-onset degradation phenomena that occur even under steady-state conditions. Continuous operation drives gradual fluoride-ion release from the membrane, catalyst-particle agglomeration (Ostwald ripening) that reduces the active electrochemical surface area, and the slow buildup of oxides on the bipolar plates, all of which progressively increase the stack's ohmic resistance [22]. By integrating these electrochemical realities into a unified techno-environmental framework, the proposed health-aware control logic achieves more than mere economic optimization. It actively suppresses the operational profiles most responsible for catalyst dissolution and membrane fatigue, thereby extending the lifecycle of critical raw materials and reinforcing the fundamental sustainability of the hydrogen production system.

7.4 Storage as a Control Resource

Hydrogen storage is often discussed as an infrastructure issue: tanks, caverns, pipelines, carriers, and safety. In this paper, storage also functions as a control resource. It allows the plant to produce hydrogen when electricity is clean and stable, rather than when demand arises. This changes the dispatch problem fundamentally.

Compressed storage is suitable for short-duration buffering, while salt caverns can be attractive for large seasonal volumes where geology permits. Derivative carriers such as ammonia, methanol, liquid hydrogen, and liquid organic hydrogen carriers may be useful for transport, but each adds conversion losses and safety requirements. The appropriate storage solution depends on scale, location, and final use. The modeling lesson is general: storage should be co-optimized with electrolyzer and renewable sizing, not added after the main design is finalized.

7.5 Water, Locality, and Environmental Realism

Water requirements are modest relative to many industrial water uses, but locality matters. A solar-rich desert site and a coastal industrial cluster face different water constraints. Freshwater withdrawals, desalinated seawater, and treated wastewater have different environmental and economic profiles. Future work should therefore add spatial water-stress indicators and water-supply pathways to the optimization model.

This extension would also improve policy relevance. Some countries are targeting hydrogen exports from regions with excellent renewable resources but limited water availability. A credible hydrogen sustainability assessment must therefore evaluate energy, carbon, storage, and water together.

Water Availability, Treatment Pathways, and Regional Constraints

The reference model accounts for water using a cost term, but project-level green hydrogen design requires a broader water boundary. Electrolysis not only requires the stoichiometric water contained in the hydrogen molecule. It also requires makeup water of sufficient purity, pretreatment, deionization or polishing, cooling, balance-of-plant water, and management of concentrate or reject streams. For PEM electrolysis, this distinction is important because the durability of the membrane, catalyst, and balance-of-plant components depends on both water quality and quantity.

$$W_{stoich,t} = 9 H_t \quad (54)$$

$$W_{req,t} = \kappa_w H_t, \kappa_w \geq 9 \quad (55)$$

here, the first expression gives the stoichiometric water requirement, while the second expression allows the model to represent system-level water demand through a water-intensity parameter. In practice, this parameter should be site-specific because it depends on purification technology, cooling configuration, water recovery, blowdown, and operating losses. The value used in the present model is therefore a simplified cost parameter rather than a complete water-supply model.

A deployment-grade version of the framework should distinguish at least three water-supply pathways: freshwater withdrawals, desalinated seawater or brackish water, and treated wastewater reuse. The corresponding water-availability constraint can be written as:

$$W_{req,t} \leq W_{fresh,t} + W_{desal,t} + W_{reuse,t} \quad (56)$$

Desalination is particularly relevant for coastal hydrogen hubs and arid regions with strong solar or wind resources. It can reduce pressure on freshwater withdrawals, but it is not environmentally neutral. Reverse osmosis or thermal desalination adds electricity or heat demand, produces brine or concentrate, and requires pretreatment, membranes, pumps, intake infrastructure, and discharge management. Kumar et al. [26] show that renewable hydrogen systems can be integrated with desalination within a techno-economic framework to enable zero-emission freshwater and electricity production. Their analysis supports the idea that water supply should be co-designed with renewable generation, hydrogen production, and storage rather than treated as a small external utility charge.

Treated wastewater reuse offers a second pathway. It may be attractive for inland industrial clusters, refineries, steel corridors, ammonia plants, and port-adjacent urban systems where municipal or industrial wastewater is available. However, wastewater reuse requires advanced treatment and quality monitoring. PEM electrolyzers are sensitive to ionic impurities, organics, chlorides, silica, and metal contaminants; therefore, reuse must include polishing to electrolyzer-grade water and account for reject streams from the treatment train. A water-reuse option can lower freshwater withdrawals, but it can also shift the burden to treatment energy, chemical use, brine disposal, and permitting.

A third option is water-energy circularity through waste-heat-assisted desalination. Electrolyzers and fuel cells reject low-grade heat that is often treated as a loss. Kumar et al. [27] provide an experimental proof of concept in which electrolyzer and fuel-cell waste heat powers desalination, enabling a more circular renewable-hydrogen configuration. This pathway is relevant because it links stack thermal management, water production, and system efficiency. It could reduce auxiliary desalination energy demand, especially in off-grid or islanded hydrogen systems, although its feasibility depends on the temperature of available waste heat, the desalination technology, heat-exchanger design, and the operating synchrony among hydrogen production, fuel-cell operation, and water demand.

Regional water stress is the main spatial constraint. A hydrogen plant in a water-abundant industrial region faces a different sustainability challenge than one in an arid export zone. A more complete model should therefore impose a site-specific freshwater withdrawal limit, for example:

$$W_{fresh,t} \leq \lambda_s W_{local,t}, 0 \leq \lambda_s \leq 1 \quad (57)$$

In this expression, the withdrawal factor represents the maximum share of locally available water that can be allocated to the hydrogen project after accounting for ecological flows, municipal needs, agriculture, and other industrial users. The constraint would prevent the optimizer from selecting a low-cost water source that is physically or socially unacceptable. It would also allow comparison among freshwater use, desalination, wastewater reuse, and waste-heat-powered water recovery under common boundary conditions.

The water-cost term in the present model can therefore be interpreted as a first-order proxy:

$$C_{water} = \sum_t (c_{fresh} W_{fresh,t} + c_{desal} W_{desal,t} + c_{reuse} W_{reuse,t}) \quad (58)$$

This proxy is sufficient for the controlled dispatch comparison developed in this paper, but it should not be read as a full water-sustainability assessment. For project siting, the dispatch model should be coupled with local hydrological data, water-quality requirements, desalination energy use, wastewater availability, brine-management constraints, and regional water-stress indicators. This extension would be especially important for large hydrogen export projects in solar-rich dry regions, where the best renewable resource may coincide with the highest water constraint.

7.6 Interpreting Environmental Results before Full Lifecycle Assessment

The environmental interpretation of the results has been adjusted to align with the model's emissions boundary. The health-aware carbon-matched strategy performs best among the simulated dispatch strategies in terms of operational electricity-related emissions, direct renewable use, curtailment, and degradation-exposure proxy. This is not the same as proving full lifecycle environmental superiority. The manuscript therefore treats the result as evidence of operational carbon and durability advantages, not as a complete sustainability ranking.

A cradle-to-grave assessment would require additional inventory terms for electrolyzer manufacturing, membrane and catalyst production, renewable-generation assets, battery buffers, hydrogen-storage equipment, compression, water treatment, stack replacement, transport infrastructure, and end-of-life management. In compact form, the broader lifecycle boundary can be expressed as:

$$CI_{H_2}^{LCA} = CI_{H_2}^{op} + CI_{H_2}^{emb} + CI_{H_2}^{water} + CI_{H_2}^{storage} + CI_{H_2}^{transport} + CI_{H_2}^{EoL} \quad (59)$$

where the operational component is the only component quantified in the present study. The remaining terms are discussed qualitatively, but not measured. This distinction matters because a strategy with very

low operational emissions could still carry non-negligible embodied burdens if it requires additional storage materials, more frequent stack replacement, energy-intensive water treatment, or transport-oriented hydrogen conversion pathways.

The conclusions have therefore been revised to avoid overstating environmental superiority. The appropriate interpretation is comparative and boundary-specific: the proposed strategy improves operational emissions and stack-conscious dispatch relative to the modeled alternatives. A full claim of lifecycle superiority would require a prospective lifecycle assessment using site-specific equipment inventories, replacement schedules, water-supply pathways, storage technologies, and end-of-life assumptions.

7.7 Technological and Control Implications

The control implication is direct: PEM electrolyzers should not merely be flexible; they should be operated with disciplined flexibility. Fast ramping is an asset, but an asset can be misused. A plant controller should distinguish between a valuable clean production block and a one-hour opportunity that creates more degradation risk than it generates in economic value.

This opens the door to digital-twin and AI-assisted extensions. The reduced-form degradation proxy used here could be replaced by a learned degradation model trained on stack voltage, current density, pressure, temperature, gas purity, membrane hydration, impurity exposure, and shutdown history. Such a model could then be used to feed into model predictive control.

A second extension is stochastic optimization. The reference model assumes known hourly profiles. Real operators face uncertainty in renewable generation, prices, demand, and grid carbon intensity. A robust or stochastic model predictive controller could update the production schedule as forecasts change, while preserving minimum up-time and degradation constraints.

7.8 Policy Implications

The policy implications are practical. First, production subsidies should not reward kilograms alone. If they do, they may favor high-utilization grid-assisted electrolysis that is cheap but not sufficiently clean. Second, certification systems should include temporal matching and transparent reporting of grid imports. Third, support schemes should recognize the system value of storage, as it enables clean and stable operating blocks. Fourth, durability should be incorporated into the hydrogen system assessment. A short-lived electrolyzer fleet is not a robust transition asset.

The broader market context makes these points urgent. The IEA [3] reports that low-emissions hydrogen remains below 1% of total production despite growth and policy attention. Odenweller and Ueckerdt [11] show that project completion has lagged announcements. If public support is limited, it should prioritize projects that can demonstrate operational credibility, including low hourly emissions, additional renewable supply, adequate storage, realistic utilization, and a credible stack lifetime.

To transition this theoretical hourly dispatch model into an operational 10 MW industrial plant, the health-aware carbon-matched algorithm must be integrated within the facility's Energy Management System (EMS) and executed via Supervisory Control and Data Acquisition (SCADA) networks [28]. While the proposed optimization provides a macro-level schedule, real-time industrial deployment requires robust computational methods—such as multivariable Model Predictive Control (MPC)—running on edge-computing Programmable Logic Controllers (PLCs) to mitigate communication latency. Execution relies strictly on high-frequency sensor telemetry and external data integration. Internally, the SCADA system must continuously monitor stack voltage, localized pressure, and cell temperature to dynamically update the degradation proxy variables (α, β, γ) and prevent thermal stress during rapid power ramps. Externally, the

EMS requires secure API connections to retrieve predictive weather data for PV/wind forecasting, day-ahead electricity pricing, and real-time grid carbon-intensity signals. Bridging the computational lag between these external signals and the microsecond response of the balance-of-plant (BoP) rectifiers is essential to executing the health-aware dispatch without inducing unintended electrochemical stress or triggering safety trip-wires in the high-pressure compressor system [29].

8 Limitations and Future Research

This study has limitations. First, the simulation uses a stylized reference system rather than measured hourly data from a specific region. That choice makes the comparison transparent, but it reduces local specificity. Future work should apply the model to industrial clusters, ports, refinery hubs, ammonia plants, and steel corridors using real renewable-generation, price, and grid-carbon data.

Implications of Using Real Site-Specific Data

The use of a calibrated reference system is useful for isolating the effect of dispatch logic, but it also means that the numerical values should not be interpreted as a site-specific forecast. With real project data, the same optimization structure would be populated with measured or reanalysis-based renewable profiles, local electricity prices, official hourly grid-carbon-intensity series, observed hydrogen offtake, local water constraints, and technology-specific stack data. The site-specific input vector can be expressed as follows:

$$X_t^{\text{site}} = (CF_{PV,t}^{\text{site}}, CF_{W,t}^{\text{site}}, p_t^{\text{site}}, CI_t^{\text{site}}, q_t^{\text{site}}, W_t^{\text{site}}) \quad (60)$$

where the elements denote, respectively, the local photovoltaic and wind capacity factors, electricity price, grid carbon intensity, hydrogen demand, and water-related input conditions. In a real deployment, these variables would replace the reference profiles used in the present simulation. Renewable availability would then be computed from the local resource mix rather than from the generic reference profile:

$$R_t^{\text{site}} = K_{PV}^{\text{site}} CF_{PV,t}^{\text{site}} + K_W^{\text{site}} CF_{W,t}^{\text{site}} \quad (61)$$

The results could change in several directions. A high-quality wind-solar site with complementary generation would increase the number of stable low-carbon production blocks, raise electrolyzer utilization, reduce curtailment, and lower LCOH. A solar-dominant site with weak night-time renewable output would make storage more valuable and could increase cycling if storage were undersized. A region with a high-carbon grid would make grid-assisted operation less attractive and would strengthen the role of carbon-admissibility constraints. Conversely, a very low-carbon grid could reduce the emissions penalty of limited grid support, although the stack-health logic would remain relevant because starts and ramps are still driven by the operating profile, not only by carbon intensity.

Site-specific demand also matters. A flat industrial offtake profile favors storage-backed production blocks, whereas a highly variable offtake profile may require larger hydrogen storage, more flexible downstream demand, or occasional low-carbon grid support. Local financing conditions and stack-service contracts would affect the annualized capital charge and the degradation-related replacement cost. Thus, the absolute values of LCOH, hydrogen output, curtailment, and operational carbon intensity are expected to be site-dependent. The comparative mechanism, however, is expected to remain relevant whenever renewable availability and hydrogen demand are temporally mismatched.

For a site-specific application, the two main reported indicators would be recomputed as:

$$CI_{H_2}^{\text{site}} = \frac{\sum_t p_{\text{grid},t}^{\text{site}} CI_t^{\text{site}}}{\sum_t H_t^{\text{site}}} \quad (62)$$

$$LCOH^{\text{site}} = \frac{C_{\text{total}}^{\text{site}}}{\sum_t H_t^{\text{site}}} \quad (63)$$

These expressions show why site-specific data can alter the numerical results. Local renewable profiles change the denominator through hydrogen output. Local grid-carbon intensity changes operational emissions. Local prices, financing conditions, stack lifetime, storage sizing, and water constraints change the total cost. The present reference case should therefore be read as a controlled comparison of operating strategies, not as a universal estimate of costs or emissions. A deployment-grade assessment would re-run the model for a defined site using measured or certified hourly data, then validate the resulting operating schedule against plant-level constraints and manufacturer-specific degradation limits.

Second, the degradation model is reduced-form. It captures starts, ramps, and operating exposure, but it does not detail electrochemical degradation mechanisms. A more rigorous engineering version should account for voltage evolution, membrane thinning, catalyst dissolution, impurity effects, pressure dynamics, and thermal cycling.

Third, the environmental boundary is operational rather than full lifecycle. The reported emissions, therefore, support conclusions about operational electricity-related performance, not about complete cradle-to-grave environmental superiority. A complete prospective lifecycle assessment should include electrolyzer manufacturing, renewable assets, storage tanks, compressors, battery buffers, water treatment, stack replacement, transport, and end-of-life management.

Fourth, the storage model is simplified. Future research should compare compressed tanks, salt caverns, pipeline line-pack, liquid hydrogen, ammonia conversion, methanol synthesis, and liquid organic carriers under common boundary conditions.

Fifth, water is represented through a simplified cost term and through a qualitative discussion of supply pathways. The model does not optimize desalination design, wastewater treatment, brine management, water-quality monitoring, or regional water-stress constraints. A deployment-grade extension should couple the dispatch model with site-specific hydrological data, desalination or reuse technology choices, and local water-allocation limits.

These limitations do not overturn the central result. They identify the next steps needed to move from reference-system simulation to deployment-grade hydrogen-system design.

9 Conclusions

This paper developed a calibrated health-aware carbon-matched dispatch framework for renewable PEM electrolysis. The model integrates hourly renewable availability, grid carbon intensity, hydrogen storage, part-load operation, degradation exposure, and cost. It then compares four operating strategies under identical system assumptions.

The results show that baseload grid-assisted electrolysis can produce the lowest apparent LCOH, but only by accepting material grid dependence and higher operational emissions. Price-only flexible operation reduces grid use but results in excessive starts. Carbon-aware flexible operation sharply reduces operational emissions, but it remains underutilized and highly fragmented. Within the operational boundary of this study, the proposed health-aware carbon-matched strategy achieves the strongest balance: very low electricity-related emissions, high direct renewable share, fewer starts, lower degradation exposure, and substantially lower LCOH than carbon-aware flexible operation. This should not be read as a full lifecycle environmental ranking.

The calibration exercise strengthens the claim. The result is not built on extreme parameter choices. It survives conservative efficiency assumptions, lower degradation penalties, and stricter carbon thresholds.

The advantage comes from a structural mechanism: clean hydrogen production should be organized into stable, low-carbon production blocks, rather than isolated clean hours.

The broader message is simple. Green hydrogen is not defined only by electrolysis technology. It is defined by how that technology is operated over time and by the lifecycle boundary used to evaluate it. A hydrogen system that ignores hourly carbon conditions may fail operational carbon tests. A hydrogen system that ignores stack health may shift the burden toward earlier replacement. Transformative hydrogen systems need carbon-matched operation, durable control, and lifecycle verification.

Acknowledgement: Not applicable.

Funding Statement: The authors received no specific funding for this study.

Author Contributions: The authors confirm contribution to the paper as follows: Conceptualization, Cosimo Magazzino and Mustafa Tevfik Kartal; methodology, Cosimo Magazzino; software, Cosimo Magazzino; validation, Cosimo Magazzino; resources, Cosimo Magazzino and Muhammad Waqas; writing—original draft preparation, Mustafa Tevfik Kartal; writing—review and editing, Cosimo Magazzino and Muhammad Waqas; visualization, Cosimo Magazzino; project administration, Cosimo Magazzino and Mustafa Tevfik Kartal; funding acquisition, Cosimo Magazzino. All authors reviewed and approved the final version of the manuscript.

Availability of Data and Materials: Data is available upon reasonable request.

Ethics Approval: Not applicable.

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

Appendix A




Figure A1: Dispatch strategy frontier for PEM electrolysis.




Figure A2: Storage-enabled smoothing of cost and cycling.




Figure A3: Economic and shutdown-penalty sensitivity landscape.




Figure A4: Robustness of the health-aware strategy under calibration variants.

References

1. Azadnia AH, McDaid C, Andwari AM, Hosseini SE. Green hydrogen supply chain risk analysis: a European hard-to-abate sectors perspective. *Renew Sustain Energy Rev.* 2023;182(C):113371. doi:10.1016/j.rser.2023.113371.
2. Franco A, Giovannini C. Routes for hydrogen introduction in the industrial hard-to-abate sectors for promoting energy transition. *Energies.* 2023;16(16):6098. doi:10.3390/en16166098.
3. IEA. Global hydrogen review 2025. Paris, France: International Energy Agency; 2025.
4. DOE. Hydrogen shot: water electrolysis technology assessment. Washington, DC, USA: U.S. Department of Energy; 2024.
5. Sayed-Ahmed H, Toldy ÁI, Santasalo-Aarnio A. Dynamic operation of proton exchange membrane electrolyzers—critical review. *Renew Sustain Energy Rev.* 2024;189(Part A):113883. doi:10.1016/j.rser.2023.113883.
6. Glenk G, Reichelstein S. Economics of converting renewable power to hydrogen. *Nat Energy.* 2019;4(3):216–22. doi:10.1038/s41560-019-0326-1.
7. Casas Ferrús MN, Ruhnau O, Madlener R. Portfolio effects in green hydrogen production under temporal matching requirements. *Energy Strategy Rev.* 2024;56(4):101580. doi:10.1016/j.esr.2024.101580.
8. Horky F, Bigerna S, Polinori P, Dimitriou D, Magazzino C. The role of financial markets in hydrogen adoption. *Eur Financ Manag.* 2026;32(1):359–82. doi:10.1111/eufm.70012.
9. Magazzino C, Haroon M. AI-based modelling and processing technologies for hydrogen creation. *J Sustain.* 2025;1(1):1–9. doi:10.55845/jos-2025-1112.
10. Magazzino C, Mele M, Leogrande A. Regional disparities and strategic implications of hydrogen production in 27 European countries. *Reg Sci Environ Econ.* 2024;1(1):3–30. doi:10.3390/rsee1010002.
11. Odenweller A, Ueckerdt F. The green hydrogen ambition and implementation gap. *Nat Energy.* 2025;10(1):110–23. doi:10.1038/s41560-024-01684-7.

12. Crespi E, Guandalini G, Mastropasqua L, Campanari S, Brouwer J. Experimental and theoretical evaluation of a 60 kW PEM electrolysis system for flexible dynamic operation. *Energy Convers Manag.* 2023;277(1):116622. doi:10.1016/j.enconman.2022.116622.
13. Wei S, Sacchi R, Tukker A, Suh S, Steubing B. Future environmental impacts of global hydrogen production. *Energy Environ Sci.* 2024;17(6):2157–72. doi:10.1039/d3ee03875k.
14. Schmidt O, Gambhir A, Staffell I, Hawkes A, Nelson J, Few S. Future cost and performance of water electrolysis: an expert elicitation study. *Int J Hydrogen Energy.* 2017;42(52):30470–92. doi:10.1016/j.ijhydene.2017.10.045.
15. Buttler A, Spliethoff H. Current status of water electrolysis for energy storage, grid balancing and sector coupling via power-to-gas and power-to-liquids: a review. *Renew Sustain Energy Rev.* 2018;82(3):2440–54. doi:10.1016/j.rser.2017.09.003.
16. DOE. Clean hydrogen production cost scenarios with PEM electrolyzer technology. In: *Hydrogen program record 24005*. Washington, DC, USA: U.S. Department of Energy; 2024.
17. Badgett A, Brauch J, Thatte A, Rubin R, Skangos C, Wang X, et al. Updated manufactured cost analysis for proton exchange membrane water electrolyzers. Golden, CO, USA: National Renewable Energy Laboratory; 2024. Report No.: NREL/TP-6A20-87625.
18. IEA. Global hydrogen review 2024. Paris, France: International Energy Agency; 2024.
19. Frensch SH, Fouda-Onana F, Serre G, Thoby D, Araya SS, Kær SK. Influence of the operation mode on PEM water electrolysis degradation. *Int J Hydrogen Energy.* 2019;44(57):29889–98. doi:10.1016/j.ijhydene.2019.09.169.
20. Weiß A, Siebel A, Bernt M, Shen TH, Tileli V, Gasteiger HA. Impact of intermittent operation on lifetime and performance of a PEM water electrolyzer. *J Electrochem Soc.* 2019;166(8):F487–97. doi:10.1149/2.0421908jes.
21. Carmo M, Fritz DL, Mergel J, Stolten D. A comprehensive review on PEM water electrolysis. *Int J Hydrogen Energy.* 2013;38(12):4901–34. doi:10.1016/j.ijhydene.2013.01.151.
22. Bernt M, Siebel A, Gasteiger HA. Analysis of voltage losses in PEM water electrolyzers with low platinum group metal loadings. *J Electrochem Soc.* 2018;165(5):F305–14. doi:10.1149/2.0641805jes.
23. Schofield L, Paren B, Macdonald R, Shao-Horn Y, Mallapragada D. Dynamic optimization of proton exchange membrane water electrolyzers considering usage-based degradation. *arXiv:2405.06766*. 2024.
24. IRENA. The potential for green hydrogen and related commodities trade. Abu Dhabi, United Arab Emirates: International Renewable Energy Agency; 2025.
25. Talukdar M, Blum P, Heinemann N, Miocic J. Techno-economic analysis of underground hydrogen storage in Europe. *iScience.* 2024;27(1):108771. doi:10.1016/j.isci.2023.108771.
26. Kumar P, Date A, Shabani B. Techno-economic analysis of an integrated desalination-renewable-hydrogen system for zero-emission freshwater and electricity production. *Energy Convers Manag.* 2026;353(5):121231. doi:10.1016/j.enconman.2026.121231.
27. Kumar P, Date A, Das RK, Shabani B. Electrolyser and fuel cell waste-heat-powered desalination enabling water–energy circularity in renewable hydrogen systems: an experimental proof of concept. *Appl Therm Eng.* 2026;293(1):130429. doi:10.1016/j.applthermaleng.2026.130429.
28. Ahaotu MC, Ogbogu CE, Thornburg J, Akwukwaegbu IO. Simulation of PEM electrolyzer power management with renewable generation in Owerri, Nigeria. *Energies.* 2025;18(1):208. doi:10.3390/en18010208.
29. Backurs A, Jansons L, Zemite L, Laizans A, Dimants J. Performance efficiency of PEM electrolyzer stacks with nano-coated titanium PTLS and bipolar plates. *Latv J Phys Tech Sci.* 2026;63(1):38–50. doi:10.2478/lpts-2026-0004.